



INVESTOR IN PEOPLE



CERTIFIED COPY OF PRIORITY DOCUMENT

The Patent Office
Concept House
Cardiff Road
Newport
South Wales
NP10 8QQ

09/866,585
Jebu Jacob Rajan
May 30, 2001

I, the undersigned, being an officer duly authorised in accordance with Section 74(1) and (4) of the Deregulation & Contracting Out Act 1994, to sign and issue certificates on behalf of the Comptroller-General, hereby certify that annexed hereto is a true copy of the documents as originally filed in connection with the patent application identified therein.

In accordance with the Patents (Companies Re-registration) Rules 1982, if a company named in this certificate and any accompanying documents has re-registered under the Companies Act 1980 with the same name as that with which it was registered immediately before re-registration save for the substitution as, or inclusion as, the last part of the name of the words "public limited company" or their equivalents in Welsh, references to the name of the company in this certificate and any accompanying documents shall be treated as references to the name with which it is so re-registered.

In accordance with the rules, the words "public limited company" may be replaced by p.l.c., plc, P.L.C. or PLC.

Re-registration under the Companies Act does not constitute a new legal entity but merely subjects the company to certain additional company law rules.


Signed

Dated

03 JUL 2001

Patents Form 7/77
Patents Act 1977
(Rule 15)

The
Patent
Office



**Statement of inventorship and of
right to grant of a patent**

1. Your reference
2722101/AM
2. Patent Application Number
accompanying application reference 2722101 17 AUG 2000 0020310.9
3. Full name of the or each applicant
Canon Kabushiki Kaisha
4. Title of the invention
SPEECH PROCESSING SYSTEM
5. State how the applicant(s) derived the right from the inventor(s) to be granted a patent
By employment of the inventor by Canon Research Centre Europe Limited, and by a general agreement dated 1 January 1994 between Canon Research Centre Europe Limited and the applicant
6. How many, if any additional Patents Forms
7/77 are attached to this form?
NONE
7. I/We believe that the person(s) named over the page (and on any extra copies of this form) is/are the inventor(s) of the invention which the above patent application relates to.
Signature Beresford & Co Date 17 August 2000
BERESFORD & Co
8. Name and daytime telephone number of
person to contact in the United Kingdom
ALAN MACDOUGALL
Tel: 020 7831 2290


Patents Form 7/77

RAJAN, Jebu Jacob
c/o Canon Research Centre Europe Limited
1 Occam Court, Occam Road
Surrey Research Park
Guildford
Surrey GU2 5YJ.

Patents Form 1/77
Patents Act 1977
(Rule 16)

The
Patent
Office

18AUG00 E561783-2 D02917
P0177108 8-00-0020310.9



Request for grant of a patent

1. Your reference
2722101/AM
2. Patent Application Number
17 AUG 2000 0020310.9
3. Full name, address and postcode of the or of each applicant (*underline all surnames*)

Canon Kabushiki Kaisha
30-2 3-Chome Shimomaruko
Ohta-Ku
Tokyo
Japan

Patents ADP number (*if known*) 00368954001

If the applicant is a corporate body, give the country/state of its incorporation Country: JAPAN
State:

4. Title of the invention
SPEECH PROCESSING SYSTEM
5. Name of agent Beresford & Co

"Address for Service" in the United Kingdom to which all correspondence should be sent 2/5 Warwick Court
High Holborn
London WC1R 5DJ

Patents ADP number 00001826001
6. Priority details

Country	Priority application number	Date of filing
UNITED KINGDOM	0013541.8	02 JUNE 2000

Patents Form 1/77

7. If this application is divided or otherwise derived from an earlier UK application give details

Number of earlier application

Date of filing

8. Is a statement of inventorship and or right to grant of a patent required in support of this request?

YES

9. Enter the number of sheets for any of the following items you are filing with this form.

Continuation sheets of this form

Description

57

Claim(s)

24

Abstract

1

Drawing(s)

11

10. If you are also filing any of the following, state how many against each item.

Priority documents

Translations of priority documents

Statement of inventorship and right to grant of a patent (*Patents form 7/77*)

1 + 1 COPY

Request for preliminary examination and search (*Patents Form 9/77*)

Request for Substantive Examination (*Patents Form 10/77*)

Any other documents (*please specify*)

11. I/We request the grant of a patent on the basis of this application

Signature

Beresford & Co
BERESFORD & Co

Date 17 August 2000

12. Name and daytime telephone number of person to contact in the United Kingdom

ALAN MACDOUGALL

Tel: 020 7831 2290

SPEECH PROCESSING SYSTEM

The present invention relates to an apparatus for and method of speech transmission. The invention particularly relates to a statistical processing of an input speech signal to derive parameter values defining the speech production system which generated the speech and the subsequent transmission of those parameter values.

A number of speech transmission systems have been proposed. These can be categorised into systems which transmit essentially the speech signals and those which parameterise the speech first (based on some speech model) and transmit the parameters. The problem with the first technique is that it requires a relatively large bandwidth to transmit the speech signals compared with parameterising the speech signals first and then transmitting the parameters. However, parameterising the speech results in more complex transmitter circuitry in order to determine the parameter values and more complex receiver circuitry in order to regenerate the speech from the transmitted parameters. Further, the second technique also reduces the quality of the regenerated speech at the receiver, since some speech information

will inevitably be lost through the parameterisation.

Many different techniques are known to parameterise speech signals. One of the most commonly used techniques is based on a linear prediction analysis of the speech. With this technique, the entire speech signal is divided into a number of time frames (typically having a duration of 10 to 30 ms) and a set of linear prediction parameters (or coefficients) is calculated to represent the speech within each time frame. This linear prediction analysis assumes that the value of a current speech sample can be predicted from a linear weighted combination of the k most recent speech samples. Based on this model, the task of the linear prediction analyses is to identify the value of the weightings (or coefficients) which minimises the mean squared error between the actual value of the current speech sample and the predicted value of the current speech sample.

One of the problems with this linear prediction analysis is that it analyses the speech within each frame in isolation from the speech within other frames. It also assumes that the same number of weightings or coefficients represent the speech within each time frame. As a result, errors are introduced into the

representation.

5 An aim of the present invention is to provide an alternative technique for parameterising speech prior to transmission from a transmitting terminal to a receiving terminal.

10 According to one aspect, the present invention provides an audio encoding system comprising: a memory for storing a probability density function for parameters of a predetermined audio model which is assumed to have generated a set of received audio signal values; means for applying the set of received audio signal values to the probability density function; means for processing
15 said function with the set of received audio signal values applied to determine samples of parameter values from said probability density function; means for analysing at least some of the determined samples to determine parameter values that are representative of the received audio signal values; and means for encoding said
20 determined parameter values to generate encoded data representative of the received audio signal values.

25 Exemplary embodiments of the present invention will now be described with reference to the accompanying drawings

in which:

5 Figure 1 is a schematic block diagram illustrating the principal components of a speech encoding and transmission system;

Figure 2a is a flow chart illustrating the processing steps performed by the transmission side of the system shown in Figure 1;

10

Figure 2b is a flow chart illustrating the processing steps performed by the receiving side of the system shown in Figure 1;

15

Figure 3 is a block diagram representing a model employed by a statistical analysis unit which forms part of the speech transmission system shown in Figure 1;

20

Figure 4 is a flow chart illustrating the processing steps performed by a model order selection unit forming part of the statistical analysis unit shown in Figure 1;

25

Figure 5 is a flow chart illustrating the main processing steps employed by a Simulation Smoother which forms part of the statistical analysis unit shown in Figure 1;

Figure 6 is a block diagram illustrating the main processing components of the statistical analysis unit shown in Figure 1;

5 Figure 7 is a memory map illustrating the data that is stored in a memory which forms part of the statistical analysis unit shown in Figure 1;

10 Figure 8 is a flow chart illustrating the main processing steps performed by the statistical analysis unit shown in Figure 6;

15 Figure 9a is a histogram for a model order of an autoregressive filter model which forms part of the model shown in Figure 3;

Figure 9b is a histogram for the variance of process noise modelled by the model shown in Figure 3; and

20 Figure 9c is a histogram for a third coefficient of the AR filter model.

25 Embodiments of the present invention can be implemented in computer hardware, computer software or a mix of computer hardware and software. When implemented using

computer software, the program instructions that make the programmable hardware operate in accordance with the present invention may be supplied on, for example, a storage device such as a magnetic disc or by downloading the software from a computer device over a computer network.

Figure 1 shows a speech encoding and transmission system. Electrical signals representative of input speech from the microphone 7 are input to a filter 15 which removes unwanted frequencies (in this embodiment frequencies above 8 kHz). The filtered signal is then sampled (at a rate of 16 kHz) and digitized by an analogue to digital converter 17 and the digitized speech samples are then stored in a buffer 19. Sequential blocks (or frames) of speech samples are then passed from the buffer 19 to a statistical analyses unit 21 which performs a statistical analysis of each frame of speech samples in sequence to determine, amongst other things, a set of auto regressive (AR) coefficients representative of the speech within the frame. In this embodiment, the AR coefficients output by the statistical analysis unit 21 are input to a channel encoder which encodes the sequences of AR filter coefficients so that they are in a more suitable form for transmission through a communications channel. The

encoded AR filter coefficients are then passed to a transmitter 73 where the encoded data is used to modulate a carrier signal which is then transmitted to a remote receiver 75. The receiver 75 demodulates the received signal to recover the encoded data which is then decoded by a decoder 76. The sequence of AR coefficients output by the decoder are then either passed to a speech recognition unit 77 which compares the sequences with stored reference models (not shown) to generate a recognition result or to a speech synthesis unit 79 which regenerates the speech and outputs it via a loudspeaker 81. As shown, prior to application to the speech synthesis unit 79, the sequences of AR coefficients may also pass through an optional processing unit 83 (shown in phantom) which can be used to manipulate the characteristics of the speech that is synthesised.

Figure 2a is a flow chart illustrating the processing steps performed by the channel encoder 71 of the system shown in Figure 1. As shown, in step S101 the channel encoder 71 receives the speech parameters and the quality indicator for the current segment of speech to be transmitted. The processing then proceeds to step S103 where the channel encoder 71 determines whether or not the speech parameters to be transmitted were generated

from a high quality speech signal. If they are, then the processing proceeds to step S105 where the channel encoder 71 encodes the speech parameters using an efficient encoding technique which minimises the data to be transmitted. If, on the other hand, step S103 determines that the speech parameters were derived from low quality speech, then the processing proceeds to step S107 where the channel encoder 71 uses a less efficient encoding technique (or no encoding) which minimises information lost in the encoding. After step S105 or S107, the processing proceeds to step S109 where the channel encoder 71 outputs the data to be transmitted to the transmitter unit 73 which transmits the encoded speech to the receiver 75. In this embodiment, the channel encoder 71 also encodes and transmits the quality indicator to the receiver 75 so that the receiver can decode the encoded speech parameters using the appropriate decoding technique. In order that the receiver can always decode the encoded speech quality indicator, the channel encoder 71 encodes this information using a standard encoding technique regardless of what the quality indicator is.

Figure 2b is a flow chart illustrating the processing steps performed by the decoder 76 shown in Figure 1. As

shown, in step S111, the decoder 76 recovers the quality indicator from the received encoded speech data. The processing then proceeds to step S113 where the decoder 76 determines whether or not the received encoded speech parameters were generated from a high quality speech signal. If they were, then the processing proceeds to step S115 where the decoder uses a decoding technique corresponding to the efficient encoding technique used in step S105. If, on the other hand, the decoder 76 determines that the received encoded speech data was generated from a low quality speech signal, then the processing proceeds to step S117 where the decoder 76 decodes the received data using a decoding technique corresponding to the less efficient encoding technique used in step S107. The processing then proceeds to step S119 where the decoded speech parameters are output either to the speech recognition unit 77 or to the speech synthesis unit 79.

20 ***Statistical Analysis Unit - Theory and Overview***

As mentioned above, the statistical analysis unit 21 analyses the speech within successive frames of the input speech signal. In most speech processing systems, the frames are overlapping. However, in this embodiment, the frames of speech are non-overlapping and have a duration

of 20ms which, with the 16kHz sampling rate of the analogue to digital converter 17, results in a frame size of 320 samples.

5 In order to perform the statistical analysis on each of the frames, the analysis unit 21 assumes that there is an underlying process which generated each sample within the frame. The model of this process used in this embodiment is shown in Figure 2. As shown, the process is modelled
10 by a speech source 31 which generates, at time $t = n$, a raw speech sample $s(n)$. Since there are physical constraints on the movement of the speech articulators, there is some correlation between neighbouring speech samples. Therefore, in this embodiment, the speech
15 source 31 is modelled by an auto regressive (AR) process. In other words, the statistical analysis unit 21 assumes that a current raw speech sample ($s(n)$) can be determined from a linear weighted combination of the most recent previous raw speech samples, i.e.:

20

$$s(n) = a_1s(n-1) + a_2s(n-2) + \dots + a_k s(n-k) + e(n) \quad (1)$$

25

where a_1, a_2, \dots, a_k are the AR filter coefficients representing the amount of correlation between the speech samples; k is the AR filter model order; and $e(n)$

represents random process noise which is involved in the generation of the raw speech samples. As those skilled in the art of speech processing will appreciate, these AR filter coefficients are the same coefficients that the linear prediction analysis estimates albeit using a different processing technique.

As shown in Figure 3, the raw speech samples $s(n)$ generated by the speech source are input to a channel 33 which models the acoustic environment between the speech source 31 and the output of the analogue to digital converter 17. Ideally, the channel 33 should simply attenuate the speech as it travels from the source 31 to the microphone. However, due to reverberation and other distortive effects, the signal $(y(n))$ output by the analogue to digital converter 17 will depend not only on the current raw speech sample $(s(n))$ but it will also depend upon previous raw speech samples. Therefore, in this embodiment, the statistical analysis unit 21 models the channel 33 by a moving average (MA) filter, i.e.:

$$y(n) = h_0 s(n) + h_1 s(n-1) + h_2 s(n-2) + \dots + h_r s(n-r) + \varepsilon(n) \quad (2)$$

where $y(n)$ represents the signal sample output by the analogue to digital converter 17 at time $t = n$; $h_0, h_1,$

$h_2 \dots h_r$ are the channel filter coefficients representing the amount of distortion within the channel 33; r is the channel filter model order; and $\varepsilon(n)$ represents a random additive measurement noise component.

5

For the current frame of speech being processed, the filter coefficients for both the speech source and the channel are assumed to be constant but unknown. Therefore, considering all N samples (where $N = 320$) in the current frame being processed gives:

10

$$\begin{aligned} s(n) &= a_1 s(n-1) + a_2 s(n-2) + \dots + a_k s(n-k) + e(n) \\ s(n-1) &= a_1 s(n-2) + a_2 s(n-3) + \dots + a_k s(n-k-1) + e(n-1) \end{aligned} \quad (3)$$

15

$$s(n-N+1) = a_1 s(n-N) + a_2 s(n-N-1) + \dots + a_k s(n-k-N+1) + e(n-N+1)$$

which can be written in vector form as:

$$s(n) = S \cdot a + e(n) \quad (4)$$

20

where

25

$$S = \begin{bmatrix} s(n-1) & s(n-2) & s(n-3) & \dots & s(n-k) \\ s(n-2) & s(n-3) & s(n-4) & \dots & s(n-k-1) \\ s(n-3) & s(n-4) & s(n-5) & \dots & s(n-k-2) \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ s(n-N) & s(n-N-1) & s(n-N-2) & \dots & s(n-k-N+1) \end{bmatrix}_{N \times k}$$

5

10

15

20

25

Similarly, considering the channel model defined by equation (2), with $h_0 = 1$ (since this provides a more stable solution), gives:

$$\begin{aligned} q(n) &= h_1 s(n-1) + h_2 s(n-2) + \dots + h_r s(n-r) + \varepsilon(n) \\ q(n-1) &= h_1 s(n-2) + h_2 s(n-3) + \dots + h_r s(n-r-1) + \varepsilon(n-1) \\ &\vdots \\ q(n-N+1) &= h_1 s(n-N) + h_2 s(n-N-1) + \dots + h_r s(n-r-N+1) + \varepsilon(n-N+1) \end{aligned} \quad (7)$$

(where $q(n) = y(n) - s(n)$) which can be written in vector form as:

$$q(n) = Y \cdot h + \varepsilon(n) \quad (8)$$

where

$$Y = \begin{bmatrix} s(n-1) & s(n-2) & s(n-3) & \dots & s(n-r) \\ s(n-2) & s(n-3) & s(n-4) & \dots & s(n-r-1) \\ s(n-3) & s(n-4) & s(n-5) & \dots & s(n-r-2) \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ s(n-N) & s(n-N-1) & s(n-N-2) & \dots & s(n-r-N+1) \end{bmatrix}_{N \times r}$$

20 and

$$h = \begin{bmatrix} h_1 \\ h_2 \\ h_3 \\ \vdots \\ h_r \end{bmatrix}_{r \times 1} \quad q(n) = \begin{bmatrix} q(n) \\ q(n-1) \\ q(n-2) \\ \vdots \\ q(n-N+1) \end{bmatrix}_{N \times 1} \quad \varepsilon(n) = \begin{bmatrix} \varepsilon(n) \\ \varepsilon(n-1) \\ \varepsilon(n-2) \\ \vdots \\ \varepsilon(n-N+1) \end{bmatrix}_{N \times 1}$$

25

In this embodiment, the analysis unit 21 aims to determine, amongst other things, values for the AR filter coefficients (a) which best represent the observed signal samples ($y(n)$) in the current frame. It does this by determining the AR filter coefficients (a) that maximise the joint probability density function of the speech model, channel model, speech samples and the noise statistics given the observed signal samples output from the analogue to digital converter 17, i.e. by determining:

$$\max_a \left\{ p(a, k, l, r, \sigma_e^2, \sigma_s^2, s(n) | y(n)) \right\} \quad (9)$$

where σ_e^2 and σ_s^2 represent the process and measurement noise statistics respectively. As those skilled in the art will appreciate, this function defines the probability that a particular speech model, channel model, raw speech samples and noise statistics generated the observed frame of speech samples ($y(n)$) from the analogue to digital converter. To do this, the statistical analysis unit 21 must determine what this function looks like. This problem can be simplified by rearranging this probability density function using Bayes law to give:

$$\frac{p(y(n)|s(n),h,r,\sigma_e^2)p(s(n)|a,k,\sigma_e^2)p(a|k)p(h|r)p(\sigma_e^2)p(\sigma_e^2)p(k)p(r)}{p(y(n))} \quad (10)$$

5 As those skilled in the art will appreciate, the denominator of equation (10) can be ignored since the probability of the signals from the analogue to digital converter is constant for all choices of model. Therefore, the AR filter coefficients that maximise the
10 function defined by equation (9) will also maximise the numerator of equation (10).

Each of the terms on the numerator of equation (10) will now be considered in turn.

15

$$p(\underline{s}(n)|\underline{a}, k, \sigma_e^2)$$

This term represents the joint probability density function for generating the vector of raw speech samples ($\underline{s}(n)$) during a frame, given the AR filter coefficients (\underline{a}), the AR filter model order (k) and the process noise statistics (σ_e^2). From equation (6) above, this joint probability density function for the raw speech samples can be determined from the joint probability density function for the process noise. In particular $p(\underline{s}(n)|\underline{a},$
20 $k, \sigma_e^2)$ is given by:

25

$$p(\underline{s}(n)|\underline{a}, k, \sigma_e^2) = p(\underline{e}(n)) \left| \frac{\delta \underline{e}(n)}{\delta \underline{s}(n)} \right|_{\underline{e}(n) = \underline{s}(n) - S\underline{a}} \quad (11)$$

5 where $p(\underline{e}(n))$ is the joint probability density function for the process noise during a frame of the input speech and the second term on the right-hand side is known as the Jacobean of the transformation. In this case, the Jacobean is unity because of the triangular form of the
10 matrix \ddot{A} (see equations (6) above).

In this embodiment, the statistical analysis unit 21 assumes that the process noise associated with the speech source 31 is Gaussian having zero mean and some unknown
15 variance σ_e^2 . The statistical analysis unit 21 also assumes that the process noise at one time point is independent of the process noise at another time point. Therefore, the joint probability density function for the process noise during a frame of the input speech (which
20 defines the probability of any given vector of process noise $\underline{e}(n)$ occurring) is given by:

$$p(\underline{e}(n)) = (2\pi\sigma_e^2)^{-\frac{N}{2}} \exp \left[\frac{-\underline{e}(n)^T \underline{e}(n)}{2\sigma_e^2} \right] \quad (12)$$

25 Therefore, the joint probability density function for a

vector of raw speech samples given the AR filter coefficients (\underline{a}), the AR filter model order (k) and the process noise variance (σ_e^2) is given by:

$$p(\underline{x}(n)|\underline{a}, k, \sigma_e^2) = (2\pi\sigma_e^2)^{-\frac{N}{2}} \exp \left[\frac{-1}{2\sigma_e^2} \left(\underline{x}(n)^T \underline{x}(n) - 2\underline{a}^T S \underline{x}(n) + \underline{a}^T S^T S \underline{a} \right) \right] \quad (13)$$

$$p(\underline{y}(n)|\underline{s}(n), \underline{h}, r, \sigma_e^2)$$

This term represents the joint probability density function for generating the vector of speech samples ($\underline{y}(n)$) output from the analogue to digital converter 17, given the vector of raw speech samples ($\underline{s}(n)$), the channel filter coefficients (\underline{h}), the channel filter model order (r) and the measurement noise statistics (σ_e^2). From equation (8), this joint probability density function can be determined from the joint probability density function for the process noise. In particular, $p(\underline{y}(n)|\underline{s}(n), \underline{h}, r, \sigma_e^2)$ is given by:

$$p(\underline{y}(n)|\underline{s}(n), \underline{h}, r, \sigma_e^2) = p(\underline{x}(n)) \left| \frac{\delta \underline{x}(n)}{\delta \underline{y}(n)} \right|_{\underline{x}(n) = \underline{q}(n) - Y \underline{h}} \quad (14)$$

where $p(\underline{x}(n))$ is the joint probability density function for the measurement noise during a frame of the input speech and the second term on the right hand side is the Jacobean of the transformation which again has a value of

one.

In this embodiment, the statistical analysis unit 21 assumes that the measurement noise is Gaussian having zero mean and some unknown variance σ_e^2 . It also assumes that the measurement noise at one time point is independent of the measurement noise at another time point. Therefore, the joint probability density function for the measurement noise in a frame of the input speech will have the same form as the process noise defined in equation (12). Therefore, the joint probability density function for a vector of speech samples ($\underline{y}(n)$) output from the analogue to digital converter 17, given the channel filter coefficients (\underline{h}), the channel filter model order (r), the measurement noise statistics (σ_e^2) and the raw speech samples ($\underline{s}(n)$) will have the following form:

$$p(\underline{y}(n)|\underline{s}(n), \underline{h}, r, \sigma_e^2) = (2\pi\sigma_e^2)^{-\frac{N}{2}} \exp \left[\frac{-1}{2\sigma_e^2} \left(\underline{q}(n)^T \underline{q}(n) - 2\underline{h}^T Y \underline{q}(n) + \underline{h}^T Y^T Y \underline{h} \right) \right] \quad (15)$$

As those skilled in the art will appreciate, although this joint probability density function for the vector of speech samples ($\underline{y}(n)$) is in terms of the variable $\underline{q}(n)$, this does not matter since $\underline{q}(n)$ is a function of $\underline{y}(n)$ and $\underline{s}(n)$, and $\underline{s}(n)$ is a given variable (ie known) for this probability density function.

$p(\underline{a}|k)$

This term defines the prior probability density function for the AR filter coefficients (\underline{a}) and it allows the statistical analysis unit 21 to introduce knowledge about what values it expects these coefficients will take. In this embodiment, the statistical analysis unit 21 models this prior probability density function by a Gaussian having an unknown variance (σ_a^2) and mean vector ($\underline{\mu}_a$), i.e.:

$$p(\underline{a}|k, \sigma_a^2, \underline{\mu}_a) = (2\pi\sigma_a^2)^{-\frac{N}{2}} \exp \left[\frac{-(\underline{a} - \underline{\mu}_a)^T (\underline{a} - \underline{\mu}_a)}{2\sigma_a^2} \right] \quad (16)$$

By introducing the new variables σ_a^2 and $\underline{\mu}_a$, the prior density functions ($p(\sigma_a^2)$ and $p(\underline{\mu}_a)$) for these variables must be added to the numerator of equation (10) above. Initially, for the first frame of speech being processed the mean vector ($\underline{\mu}_a$) can be set to zero and for the second and subsequent frames of speech being processed, it can be set to the mean vector obtained during the processing of the previous frame. In this case, $p(\underline{\mu}_a)$ is just a Dirac delta function located at the current value of $\underline{\mu}_a$ and can therefore be ignored.

With regard to the prior probability density function for

the variance of the AR filter coefficients, the statistical analysis unit 21 could set this equal to some constant to imply that all variances are equally probable. However, this term can be used to introduce knowledge about what the variance of the AR filter coefficients is expected to be. In this embodiment, since variances are always positive, the statistical analysis unit 21 models this variance prior probability density function by an Inverse Gamma function having parameters α_a and β_a , i.e.:

$$p(\sigma_a^2 | \alpha_a, \beta_a) = \frac{(\sigma_a^2)^{-(\alpha_a + 1)}}{\beta_a \Gamma(\alpha_a)} \exp \left[\frac{-1}{\sigma_a^2 \beta_a} \right] \quad (17)$$

At the beginning of the speech being processed, the statistical analysis unit 21 will not have much knowledge about the variance of the AR filter coefficients. Therefore, initially, the statistical analysis unit 21 sets the variance σ_a^2 and the α and β parameters of the Inverse Gamma function to ensure that this probability density function is fairly flat and therefore non-informative. However, after the first frame of speech has been processed, these parameters can be set more accurately during the processing of the next frame of speech by using the parameter values calculated during the processing of the previous frame of speech.

$p(\underline{h}|r)$

This term represents the prior probability density function for the channel model coefficients (\underline{h}) and it allows the statistical analysis unit 21 to introduce knowledge about what values it expects these coefficients to take. As with the prior probability density function for the AR filter coefficients, in this embodiment, this probability density function is modelled by a Gaussian having an unknown variance (σ_h^2) and mean vector ($\underline{\mu}_h$), i.e.:

$$p(\underline{h}|r, \sigma_h^2, \underline{\mu}_h) = (2\pi\sigma_h^2)^{-\frac{N}{2}} \exp\left[\frac{-(\underline{h}-\underline{\mu}_h)^T(\underline{h}-\underline{\mu}_h)}{2\sigma_h^2}\right] \quad (18)$$

Again, by introducing these new variables, the prior density functions ($p(\sigma_h)$ and $p(\underline{\mu}_h)$) must be added to the numerator of equation (10). Again, the mean vector can initially be set to zero and after the first frame of speech has been processed and for all subsequent frames of speech being processed, the mean vector can be set to equal the mean vector obtained during the processing of the previous frame. Therefore, $p(\underline{\mu}_h)$ is also just a Dirac delta function located at the current value of $\underline{\mu}_h$ and can be ignored.

With regard to the *prior* probability density function for the variance of the channel filter coefficients, again, in this embodiment, this is modelled by an Inverse Gamma function having parameters α_h and β_h . Again, the variance (σ_h^2) and the α and β parameters of the Inverse Gamma function can be chosen initially so that these densities are non-informative so that they will have little effect on the subsequent processing of the initial frame.

10 $p(\sigma_e^2)$ and $p(\sigma_\epsilon^2)$

These terms are the *prior* probability density functions for the process and measurement noise variances and again, these allow the statistical analysis unit 21 to introduce knowledge about what values it expects these noise variances will take. As with the other variances, in this embodiment, the statistical analysis unit 21 models these by an Inverse Gamma function having parameters α_e , β_e and α_ϵ , β_ϵ respectively. Again, these variances and these Gamma function parameters can be set initially so that they are non-informative and will not appreciably affect the subsequent calculations for the initial frame.

25 $p(k)$ and $p(r)$

These terms are the *prior* probability density functions

for the AR filter model order (k) and the channel model order (r) respectively. In this embodiment, these are modelled by a uniform distribution up to some maximum order. In this way, there is no prior bias on the number of coefficients in the models except that they can not exceed these predefined maximums. In this embodiment, the maximum AR filter model order (k) is thirty and the maximum channel model order (r) is one hundred and fifty.

Therefore, inserting the relevant equations into the numerator of equation (10) gives the following joint probability density function which is proportional to $p(\underline{a}, k, \underline{h}, r, \sigma_a^2, \sigma_h^2, \sigma_e^2, \sigma_c^2, \underline{s}(n) | \underline{y}(n))$:

$$\begin{aligned}
 & (2\pi\sigma_e^2)^{-\frac{N}{2}} \exp \left[\frac{-1}{2\sigma_e^2} (\underline{q}(n)^T \underline{q}(n) - 2\mathbf{h}^T \mathbf{Y} \underline{q}(n) + \mathbf{h}^T \mathbf{Y}^T \mathbf{Y} \mathbf{h}) \right] \\
 & \times (2\pi\sigma_c^2)^{-\frac{N}{2}} \exp \left[\frac{-1}{2\sigma_c^2} (\underline{s}(n)^T \underline{s}(n) - 2\mathbf{a}^T \mathbf{S} \underline{s}(n) + \mathbf{a}^T \mathbf{S}^T \mathbf{S} \mathbf{a}) \right] \\
 & \times (2\pi\sigma_a^2)^{-\frac{N}{2}} \exp \left[\frac{-(\mathbf{a} - \underline{\mu}_a)^T (\mathbf{a} - \underline{\mu}_a)}{2\sigma_a^2} \right] \times (2\pi\sigma_h^2)^{-\frac{N}{2}} \exp \left[\frac{-(\mathbf{h} - \underline{\mu}_h)^T (\mathbf{h} - \underline{\mu}_h)}{2\sigma_h^2} \right] \\
 & \times \frac{(\sigma_a^2)^{-(\alpha_a+1)}}{\beta_a \Gamma(\alpha_a)} \exp \left[\frac{-1}{\sigma_a^2 \beta_a} \right] \times \frac{(\sigma_h^2)^{-(\alpha_h+1)}}{\beta_h \Gamma(\alpha_h)} \exp \left[\frac{-1}{\sigma_h^2 \beta_h} \right] \\
 & \times \frac{(\sigma_e^2)^{-(\alpha_e+1)}}{\beta_e \Gamma(\alpha_e)} \exp \left[\frac{-1}{\sigma_e^2 \beta_e} \right] \times \frac{(\sigma_c^2)^{-(\alpha_c+1)}}{\beta_c \Gamma(\alpha_c)} \exp \left[\frac{-1}{\sigma_c^2 \beta_c} \right]
 \end{aligned}$$

(19)

Gibbs Sampler

In order to determine the form of this joint probability density function, the statistical analysis unit 21 "draws samples" from it. In this embodiment, since the joint probability density function to be sampled is a complex multivariate function, a Gibbs sampler is used which breaks down the problem into one of drawing samples from probability density functions of smaller dimensionality. In particular, the Gibbs sampler proceeds by drawing random variates from conditional densities as follows:

first iteration

$$p(a, k | h^0, r^0, \sigma_e^{20}, \sigma_\epsilon^{20}, \sigma_a^{20}, \sigma_h^{20}, s(n)^0, y(n)) \rightarrow a^1, k^1$$

$$p(h, r | a^1, k^1, \sigma_e^{20}, \sigma_\epsilon^{20}, \sigma_a^{20}, \sigma_h^{20}, s(n)^0, y(n)) \rightarrow h^1, k^1$$

$$p(\sigma_e^2 | a^1, k^1, h^1, r^1, \sigma_\epsilon^{20}, \sigma_a^{20}, \sigma_h^{20}, s(n)^0, y(n)) \rightarrow \sigma_e^{21}$$

$$p(\sigma_h^{21} | a^1, k^1, h^1, r^1, \sigma_e^{21}, \sigma_\epsilon^{21}, \sigma_a^{21}, \sigma_h^{21}, s(n)^0, y(n)) \rightarrow \sigma_h^{21}$$

second iteration

$$p(a, k | h^1, r^1, \sigma_e^{21}, \sigma_\epsilon^{21}, \sigma_a^{21}, \sigma_h^{21}, s(n)^1, y(n)) \rightarrow a^2, k^2$$

$$p(h, r | a^2, k^2, \sigma_e^{21}, \sigma_\epsilon^{21}, \sigma_a^{21}, \sigma_h^{21}, s(n)^1, y(n)) \rightarrow h^2, r^2$$

etc.

where $(h^0, r^0, (\sigma_e^2)^0, (\sigma_\epsilon^2)^0, (\sigma_a^2)^0, (\sigma_h^2)^0, \underline{s}(n)^0)$ are initial values which may be obtained from the results of the statistical analysis of the previous frame of speech, or where there are no previous frames, can be set to appropriate values that will be known to those skilled in the art of speech processing.

As those skilled in the art will appreciate, these conditional densities are obtained by inserting the current values for the given (or known) variables into the terms of the density function of equation (19). For the conditional density $p(\underline{a}, k | \dots)$ this results in:

$$p(\underline{a}, k | \dots) \propto \exp \left[\frac{-1}{2\sigma_e^2} \left(\underline{s}(n)^T \underline{s}(n) - 2\underline{a}^T S \underline{s}(n) + \underline{a}^T S^T S \underline{a} \right) \right] \quad (20)$$

$$\times \exp \left[\frac{-(\underline{a} - \underline{\mu}_a)^T (\underline{a} - \underline{\mu}_a)}{2\sigma_a^2} \right]$$

which can be simplified to give:

$$p(\underline{a}, k | \dots) \propto \exp \left[\frac{-1}{2} \left(\frac{\underline{s}(n)^T \underline{s}(n)}{\sigma_e^2} + \frac{\underline{\mu}_a^T \underline{\mu}_a}{\sigma_a^2} - 2\underline{a}^T \left[\frac{S \underline{s}(n)}{\sigma_e^2} + \frac{\underline{\mu}_a}{\sigma_a^2} \right] + \underline{a}^T \left[\frac{S^T S}{\sigma_e^2} + \frac{I}{\sigma_a^2} \right] \underline{a} \right) \right] \quad (21)$$

which is in the form of a standard Gaussian distribution having the following covariance matrix:

$$\Sigma_a = \left[\frac{S^T S}{\sigma_e^2} + \frac{I}{\sigma_a^2} \right]^{-1} \quad (22)$$

The mean value of this Gaussian distribution can be determined by differentiating the exponent of equation (21) with respect to \underline{a} and determining the value of \underline{a} which makes the differential of the exponent equal to zero. This yields a mean value of:

$$\hat{\underline{\mu}}_a = \left[\frac{S^T S}{\sigma_e^2} + \frac{I}{\sigma_a^2} \right]^{-1} \left[\frac{S^T \underline{s}(n)}{\sigma_e^2} + \frac{\underline{\mu}_a}{\sigma_a^2} \right] \quad (23)$$

A sample can then be drawn from this standard Gaussian distribution to give \underline{a}^g (where g is the g^{th} iteration of the Gibbs sampler) with the model order (k^g) being determined by a model order selection routine which will be described later. The drawing of a sample from this Gaussian distribution may be done by using a random number generator which generates a vector of random values which are uniformly distributed and then using a transformation of random variables using the covariance matrix and the mean value given in equations (22) and (23) to generate the sample. In this embodiment, however, a random number generator is used which generates random numbers from a Gaussian distribution having zero mean and a variance of one. This simplifies

the transformation process to one of a simple scaling using the covariance matrix given in equation (22) and shifting using the mean value given in equation (23). Since the techniques for drawing samples from Gaussian distributions are well known in the art of statistical analysis, a further description of them will not be given here. A more detailed description and explanation can be found in the book entitled "Numerical Recipes in C", by W. Press et al, Cambridge University Press, 1992 and in particular at chapter 7.

As those skilled in the art will appreciate, however, before a sample can be drawn from this Gaussian distribution, estimates of the raw speech samples must be available so that the matrix S and the vector $\underline{s}(n)$ are known. The way in which these estimates of the raw speech samples are obtained in this embodiment will be described later.

A similar analysis for the conditional density $p(h, r | \dots)$ reveals that it also is a standard Gaussian distribution but having a covariance matrix and mean value given by:

$$\Sigma_h = \left[\frac{Y^T Y}{\sigma_e^2} + \frac{I}{\sigma_h^2} \right]^{-1} \quad \hat{\mu}_h = \left[\frac{Y^T Y}{\sigma_e^2} + \frac{I}{\sigma_h^2} \right]^{-1} \left[\frac{Y^T q(n)}{\sigma_e^2} + \frac{\mu_h}{\sigma_h^2} \right] \quad (24)$$

from which a sample for h^g can be drawn in the manner described above, with the channel model order (r^g) being determined using the model order selection routine which will be described later.

5

A similar analysis for the conditional density $p(\sigma_e^2 | \dots)$ shows that:

10

$$p(\sigma_e^2 | \dots) \propto (\sigma_e^2)^{-\frac{N}{2}} \exp\left[\frac{-E}{2\sigma_e^2}\right] \frac{(\sigma_e^2)^{-(\alpha_e+1)}}{\beta_e \Gamma(\alpha_e)} \exp\left[\frac{-1}{\sigma_e^2 \beta_e}\right] \quad (25)$$

where:

15

$$E = \mathbf{x}(n)^T \mathbf{x}(n) - 2\mathbf{a}^T S \mathbf{x}(n) + \mathbf{a}^T S^T S \mathbf{a}$$

which can be simplified to give:

$$p(\sigma_e^2 | \dots) \propto (\sigma_e^2)^{-\left[\left(\frac{N}{2} + \alpha_e\right) + 1\right]} \exp\left[\frac{-1}{\sigma_e^2} \left(\frac{E}{2} + \frac{1}{\beta_e}\right)\right] \quad (26)$$

20

which is also an Inverse Gamma distribution having the following parameters:

$$\hat{\alpha}_e = \frac{N}{2} + \alpha_e \quad \text{and} \quad \hat{\beta}_e = \frac{2\beta_e}{2 + \beta_e E} \quad (27)$$

25

5 A sample is then drawn from this Inverse Gamma distribution by firstly generating a random number from a uniform distribution and then performing a transformation of random variables using the alpha and beta parameters given in equation (27), to give $(\sigma_e^2)^g$.

A similar analysis for the conditional density $p(\sigma_e^2|\dots)$ reveals that it also is an Inverse Gamma distribution having the following parameters:

10

$$\hat{\alpha}_e = \frac{N}{2} + \alpha_e \quad \text{and} \quad \hat{\beta}_e = \frac{2\beta_e}{2 + \beta_e E^*} \quad (28)$$

where:

$$E^* = q(n)^T q(n) - 2h^T Y q(n) + h^T Y^T Y h$$

15

A sample is then drawn from this Inverse Gamma distribution in the manner described above to give $(\sigma_e^2)^g$.

20 A similar analysis for conditional density $p(\sigma_a^2|\dots)$ reveals that it too is an Inverse Gamma distribution having the following parameters:

$$\hat{\alpha}_a = \frac{N}{2} + \alpha_a \quad \text{and} \quad \hat{\beta}_a = \frac{2\beta_a}{2 + \beta_a (a - \mu_a)^T (a - \mu_a)} \quad (29)$$

25 A sample is then drawn from this Inverse Gamma

distribution in the manner described above to give $(\sigma_a^2)^q$.

Similarly, the conditional density $p(\sigma_h^2|\dots)$ is also an Inverse Gamma distribution but having the following parameters:

$$\hat{\alpha}_h = \frac{N}{2} + \alpha_h \quad \text{and} \quad \hat{\beta}_h = \frac{2\beta_h}{2 + \beta_h(h - \mu_h)^T(h - \mu_h)} \quad (30)$$

A sample is then drawn from this Inverse Gamma distribution in the manner described above to give $(\sigma_h^2)^q$.

As those skilled in the art will appreciate, the Gibbs sampler requires an initial transient period to converge to equilibrium (known as burn-in). Eventually, after L iterations, the sample $(\underline{a}^L, k^L, \underline{h}^L, r^L, (\sigma_a^2)^L, (\sigma_e^2)^L, (\sigma_s^2)^L, (\sigma_h^2)^L, s(n)^L)$ is considered to be a sample from the joint probability density function defined in equation (19). In this embodiment, the Gibbs sampler performs approximately one hundred and fifty (150) iterations on each frame of input speech and discards the samples from the first fifty iterations and uses the rest to give a picture (a set of histograms) of what the joint probability density function defined in equation (19) looks like. From these histograms, the set of AR coefficients (\underline{a}) which best represents the observed

speech samples ($y(n)$) from the analogue to digital converter 17 are determined. The histograms are also used to determine appropriate values for the variances and channel model coefficients (\underline{h}) which can be used as the initial values for the Gibbs sampler when it processes the next frame of speech.

Model Order Selection

As mentioned above, during the Gibbs iterations, the model order (k) of the AR filter and the model order (r) of the channel filter are updated using a model order selection routine. In this embodiment, this is performed using a technique called "Reversible jump Markov chain Monte Carlo computation". This technique is described in the paper entitled "Reversible jump Markov chain Monte Carlo Computation and Bayesian model determination" by Peter Green, *Biometrika*, vol 82, pp 711 to 732, 1995.

Figure 4 is a flow chart which illustrates the processing steps performed during this model order selection routine for the AR filter model order (k). As shown, in step s1, a new model order (k_2) is proposed. In this embodiment, the new model order will normally be proposed as $k_2 = k_1 \pm 1$, but occasionally it will be proposed as $k_2 = k_1 \pm 2$.

and very occasionally as $k_2 = k_1 \pm 3$ etc. To achieve this, a sample is drawn from a discretised Laplacian density function centred on the current model order (k_1) and with the variance of this Laplacian density function being chosen *a priori* in accordance with the degree of sampling of the model order space that is required.

The processing then proceeds to step s3 where a model order variable (MO) is set equal to:

$$MO = \max \left\{ \frac{p(\underline{a}_{<1:k_2>}, k_2 | \dots)}{p(\underline{a}_{<1:k_1>}, k_1 | \dots)}, 1 \right\} \quad (31)$$

where the ratio term is the ratio of the conditional probability given in equation (21) evaluated for the current AR filter coefficients (\underline{a}) drawn by the Gibbs sampler for the current model order (k_1) and for the proposed new model order (k_2). If $k_2 > k_1$, then the matrix S must first be resized and then a new sample must be drawn from the Gaussian distribution having the mean vector and covariance matrix defined by equations (22) and (23) (determined for the resized matrix S), to provide the AR filter coefficients ($\underline{a}_{<1:k_2>}$) for the new model order (k_2). If $k_2 < k_1$ then all that is required is to delete the last ($k_1 - k_2$) samples of the \underline{a} vector. If the ratio in equation (31) is greater than one, then this

implies that the proposed model order (k_2) is better than the current model order whereas if it is less than one then this implies that the current model order is better than the proposed model order. However, since
5 occasionally this will not be the case, rather than deciding whether or not to accept the proposed model order by comparing the model order variable (MO) with a fixed threshold of one, in this embodiment, the model order variable (MO) is compared, in step s5, with a
10 random number which lies between zero and one. If the model order variable (MO) is greater than this random number, then the processing proceeds to step s7 where the model order is set to the proposed model order (k_2) and a count associated with the value of k_2 is incremented.
15 If, on the other hand, the model order variable (MO) is smaller than the random number, then the processing proceeds to step s9 where the current model order is maintained and a count associated with the value of the current model order (k_1) is incremented. The processing
20 then ends.

This model order selection routine is carried out for both the model order of the AR filter model and for the model order of the channel filter model. This routine
25 may be carried out at each Gibbs iteration. However,

this is not essential. Therefore, in this embodiment, this model order updating routine is only carried out every third Gibbs iteration.

5 ***Simulation Smoother***

As mentioned above, in order to be able to draw samples using the Gibbs sampler, estimates of the raw speech samples are required to generate $\underline{s}(n)$, S and Y which are used in the Gibbs calculations. These could be obtained
10 from the conditional probability density function $p(\underline{s}(n)|\dots)$. However, this is not done in this embodiment because of the high dimensionality of $\underline{s}(n)$. Therefore, in this embodiment, a different technique is used to provide the necessary estimates of the raw speech
15 samples. In particular, in this embodiment, a "Simulation Smoother" is used to provide these estimates. This Simulation Smoother was proposed by Piet de Jong in the paper entitled "The Simulation Smoother for Time Series Models", *Biometrika* (1995), vol 82,2, pages 339 to
20 350. As those skilled in the art will appreciate, the Simulation Smoother is run before the Gibbs Sampler. It is also run again during the Gibbs iterations in order to update the estimates of the raw speech samples. In this embodiment, the Simulation Smoother is run every fourth
25 Gibbs iteration.

In order to run the Simulation Smoother, the model equations defined above in equations (4) and (6) must be written in "state space" format as follows:

$$\begin{aligned}\hat{s}(n) &= \tilde{A} \cdot \hat{s}(n-1) + \hat{e}(n) \\ y(n) &= h^T \cdot \hat{s}(n-1) + \varepsilon(n)\end{aligned}\tag{32}$$

where

$$\tilde{A} = \begin{bmatrix} a_1 & a_2 & a_3 & \dots & a_k & 0 & \dots & 0 \\ 1 & 0 & 0 & \dots & 0 & 0 & \dots & 0 \\ 0 & 1 & 0 & \dots & 0 & 0 & \dots & 0 \\ \vdots & & & & & & & \\ \vdots & & & & & & & \\ \vdots & & & & & & & \\ 0 & & & & & & 1 & 0 \end{bmatrix}_{rxr}$$

and

$$\hat{s}(n) = \begin{bmatrix} \hat{s}(n) \\ \hat{s}(n-1) \\ \hat{s}(n-2) \\ \vdots \\ \vdots \\ \hat{s}(n-r+1) \end{bmatrix}_{rx1} \quad \hat{e}(n) = \begin{bmatrix} \hat{e}(n) \\ 0 \\ 0 \\ \vdots \\ \vdots \\ 0 \end{bmatrix}_{rx1}$$

With this state space representation, the dimensionality of the raw speech vectors ($\hat{s}(n)$) and the process noise vectors ($\hat{e}(n)$) do not need to be $N \times 1$ but only have to be as large as the greater of the model orders - k and r . Typically, the channel model order (r) will be larger than the AR filter model order (k). Hence, the vector of

raw speech samples ($\hat{s}(n)$) and the vector of process noise ($\hat{e}(n)$) only need to be rx1 and hence the dimensionality of the matrix \bar{A} only needs to be rxr.

5 The Simulation Smoother involves two stages - a first stage in which a Kalman filter is run on the speech samples in the current frame and then a second stage in which a "smoothing" filter is run on the speech samples in the current frame using data obtained from the Kalman filter stage. Figure 5 is a flow chart illustrating the processing steps performed by the Simulation Smoother. As shown, in step s21, the system initialises a time variable t to equal one. During the Kalman filter stage, this time variable is run from $t = 1$ to N in order to process the N speech samples in the current frame being processed in time sequential order. After step s21, the processing then proceeds to step s23, where the following Kalman filter equations are computed for the current speech sample ($y(t)$) being processed:

20

$$w(t) = y(t) - h^T \hat{s}(t)$$

$$d(t) = h^T P(t) h + \sigma_e^2$$

$$k_f(t) = (\tilde{A} P(t) h) . d(t)^{-1}$$

$$\hat{s}(t+1) = \tilde{A} \hat{s}(t) + k_f(t) . w(t)$$

(33)

25

$$L(t) = \tilde{A} - k_f(t) . h^T$$

$$P(t+1) = \tilde{A} P(t) L(t)^T + \sigma_e^2 . I$$

where the initial vector of raw speech samples ($\hat{s}(1)$) includes raw speech samples obtained from the processing of the previous frame (or if there are no previous frames then $s(i)$ is set equal to zero for $i < 1$); $P(1)$ is the variance of $\hat{s}(1)$ (which can be obtained from the previous frame or initially can be set to σ_e^2); h is the current set of channel model coefficients which can be obtained from the processing of the previous frame (or if there are no previous frames then the elements of h can be set to their expected values - zero); $y(t)$ is the current speech sample of the current frame being processed and I is the identity matrix. The processing then proceeds to step s25 where the scalar values $w(t)$ and $d(t)$ are stored together with the rxr matrix $L(t)$ (or alternatively the Kalman filter gain vector $k_f(t)$ could be stored from which $L(t)$ can be generated). The processing then proceeds to step s27 where the system determines whether or not all the speech samples in the current frame have been processed. If they have not, then the processing proceeds to step s29 where the time variable t is incremented by one so that the next sample in the current frame will be processed in the same way. Once all N samples in the current frame have been processed in this way and the corresponding values stored, the first stage of the Simulation Smoother is complete.

The processing then proceeds to step s31 where the second stage of the Simulation Smoother is started in which the smoothing filter processes the speech samples in the current frame in reverse sequential order. As shown, in step s31 the system runs the following set of smoothing filter equations on the current speech sample being processed together with the stored Kalman filter variables computed for the current speech sample being processed:

$$C(t) = \sigma_e^2 (I - \sigma_e^2 U(t))$$

$$\eta(t) \sim N(0, C(t))$$

$$V(t) = \sigma_e^2 U(t) L(t)$$

$$\underline{r}(t-1) = h d(t)^{-1} w(t) + L(t)^T \underline{r}(t) - V(t)^T C(t)^{-1} \eta(t)$$

$$U(t-1) = h d(t)^{-1} h^T + L(t)^T U(t) L(t) + V(t)^T C(t)^{-1} V(t)$$

(34)

$$\tilde{\underline{e}}(t) = \sigma_e^2 \underline{r}(t) + \eta(t) \quad \text{where } \tilde{\underline{e}}(t) = [\tilde{e}(t) \tilde{e}(t-1) \tilde{e}(t-2) \dots \tilde{e}(t-r+1)]^T$$

$$\hat{\underline{s}}(t) = \tilde{A} \hat{\underline{s}}(t-1) + \tilde{\underline{e}}(t) \quad \text{where } \hat{\underline{s}}(t) = [\hat{s}(t) \hat{s}(t-1) \hat{s}(t-2) \dots \hat{s}(t-r+1)]^T$$

$$\text{and } \hat{\underline{e}}(t) = [\tilde{\underline{e}}(t) \quad 0 \quad 0 \quad \dots \quad 0]^T$$

where $\eta(t)$ is a sample drawn from a Gaussian distribution having zero mean and covariance matrix $C(t)$; the initial vector $\underline{r}(t=N)$ and the initial matrix $U(t=N)$ are both set to zero; and $\underline{s}(0)$ is obtained from the processing of the previous frame (or if there are no previous frames can be set equal to zero). The processing then proceeds to step s33 where the estimate of the process noise ($\tilde{e}(t)$) for

the current speech sample being processed and the estimate of the raw speech sample ($\hat{s}(t)$) for the current speech sample being processed are stored. The processing then proceeds to step s35 where the system determines whether or not all the speech samples in the current frame have been processed. If they have not, then the processing proceeds to step s37 where the time variable t is decremented by one so that the previous sample in the current frame will be processed in the same way. Once all N samples in the current frame have been processed in this way and the corresponding process noise and raw speech samples have been stored, the second stage of the Simulation Smoother is complete and an estimate of $\underline{s}(n)$ will have been generated.

As shown in equations (4) and (8), the matrix S and the matrix Y require raw speech samples $s(n-N-1)$ to $s(n-N-k+1)$ and $s(n-N-1)$ to $s(n-N-r+1)$ respectively in addition to those in $\underline{s}(n)$. These additional raw speech samples can be obtained either from the processing of the previous frame of speech or if there are no previous frames, they can be set to zero. With these estimates of raw speech samples, the Gibbs sampler can be run to draw samples from the above described probability density functions.

Statistical Analysis Unit - Operation

A description has been given above of the theory underlying the statistical analysis unit 21. A description will now be given with reference to Figures 6 to 8 of the operation of the statistical analysis unit 21 that is used in the embodiment.

Figure 6 is a block diagram illustrating the principal components of the statistical analysis unit 21 of this embodiment. As shown, it comprises the above described Gibbs sampler 41, Simulation Smoother 43 (including the Kalman filter 43-1 and smoothing filter 43-2) and model order selector 45. It also comprises a memory 47 which receives the speech samples of the current frame to be processed, a data analysis unit 49 which processes the data generated by the Gibbs sampler 41 and the model order selector 45 and a controller 50 which controls the operation of the statistical analysis unit 21.

As shown in Figure 6, the memory 47 includes a non volatile memory area 47-1 and a working memory area 47-2. The non volatile memory 47-1 is used to store the joint probability density function given in equation (19) above and the equations for the variances and mean values and the equations for the Inverse Gamma parameters given

above in equations (22) to (24) and (27) to (30) for the above mentioned conditional probability density functions for use by the Gibbs sampler 41. The non volatile memory 47-1 also stores the Kalman filter equations given above in equation (33) and the smoothing filter equations given above in equation 34 for use by the Simulation Smoother 43.

Figure 7 is a schematic diagram illustrating the parameter values that are stored in the working memory area (RAM) 47-2. As shown, the RAM includes a store 51 for storing the speech samples $y_f(1)$ to $y_f(N)$ output by the analogue to digital converter 17 for the current frame (f) being processed. As mentioned above, these speech samples are used in both the Gibbs sampler 41 and the Simulation Smoother 43. The RAM 47-2 also includes a store 53 for storing the initial estimates of the model parameters ($g=0$) and the M samples ($g = 1$ to M) of each parameter drawn from the above described conditional probability density functions by the Gibbs sampler 41 for the current frame being processed. As mentioned above, in this embodiment, M is 100 since the Gibbs sampler 41 performs 150 iterations on each frame of input speech with the first fifty samples being discarded. The RAM 47-2 also includes a store 55 for storing $W(t)$, $d(t)$ and

$L(t)$ for $t = 1$ to N which are calculated during the processing of the speech samples in the current frame of speech by the above described Kalman filter 43-1. The RAM 47-2 also includes a store 57 for storing the estimates of the raw speech samples ($\hat{s}_f(t)$) and the estimates of the process noise ($\tilde{e}_f(t)$) generated by the smoothing filter 43-2, as discussed above. The RAM 47-2 also includes a store 59 for storing the model order counts which are generated by the model order selector 45 when the model orders for the AR filter model and the channel model are updated.

Figure 8 is a flow diagram illustrating the control program used by the controller 50, in this embodiment, to control the processing operations of the statistical analysis unit 21. As shown, in step s41, the controller 50 retrieves the next frame of speech samples to be processed from the buffer 19 and stores them in the memory store 51. The processing then proceeds to step s43 where initial estimates for the channel model, raw speech samples and the process noise and measurement noise statistics are set and stored in the store 53. These initial estimates are either set to be the values obtained during the processing of the previous frame of speech or, where there are no previous frames of speech,

are set to their expected values (which may be zero). The processing then proceeds to step s45 where the Simulation Smoother 43 is activated so as to provide an estimate of the raw speech samples in the manner described above. The processing then proceeds to step s47 where one iteration of the Gibbs sampler 41 is run in order to update the channel model, speech model and the process and measurement noise statistics using the raw speech samples obtained in step s45. These updated parameter values are then stored in the memory store 53.

The processing then proceeds to step s49 where the controller 50 determines whether or not to update the model orders of the AR filter model and the channel model. As mentioned above, in this embodiment, these model orders are updated every third Gibbs iteration. If the model orders are to be updated, then the processing proceeds to step s51 where the model order selector 45 is used to update the model orders of the AR filter model and the channel model in the manner described above. If at step s49 the controller 50 determines that the model orders are not to be updated, then the processing skips step s51 and the processing proceeds to step s53. At step s53, the controller 50 determines whether or not to perform another Gibbs iteration. If another iteration is

to be performed, then the processing proceeds to decision block s55 where the controller 50 decides whether or not to update the estimates of the raw speech samples ($s(t)$). If the raw speech samples are not to be updated, then the processing returns to step s47 where the next Gibbs iteration is run.

As mentioned above, in this embodiment, the Simulation Smoother 43 is run every fourth Gibbs iteration in order to update the raw speech samples. Therefore, if the controller 50 determines, in step s55 that there has been four Gibbs iterations since the last time the speech samples were updated, then the processing returns to step s45 where the Simulation Smoother is run again to provide new estimates of the raw speech samples ($s(t)$). Once the controller 50 has determined that the required 150 Gibbs iterations have been performed, the controller 50 causes the processing to proceed to step s57 where the data analysis unit 49 analyses the model order counts generated by the model order selector 45 to determine the model orders for the AR filter model and the channel model which best represents the current frame of speech being processed. The processing then proceeds to step s59 where the data analysis unit 49 analyses the samples drawn from the conditional densities by the Gibbs sampler

41 to determine the AR filter coefficients (a), the channel model coefficients (h), the variances of these coefficients and the process and measurement noise variances which best represent the current frame of speech being processed. The processing then proceeds to step s61 where the controller 50 determines whether or not there is any further speech to be processed. If there is more speech to be processed, then processing returns to step S41 and the above process is repeated for the next frame of speech. Once all the speech has been processed in this way, the processing ends.

Data Analysis unit

A more detailed description of the data analysis unit 49 will now be given with reference to Figure 9. As mentioned above, the data analysis unit 49 initially determines, in step s57, the model orders for both the AR filter model and the channel model which best represents the current frame of speech being processed. It does this using the counts that have been generated by the model order selector 45 when it was run in step s51. These counts are stored in the store 59 of the RAM 47-2. In this embodiment, in determining the best model orders, the data analysis unit 49 identifies the model order having the highest count. Figure 9a is an exemplary

histogram which illustrates the distribution of counts that is generated for the model order (k) of the AR filter model. Therefore, in this example, the data analysis unit 49 would set the best model order of the AR filter model as five. The data analysis unit 49 performs a similar analysis of the counts generated for the model order (r) of the channel model to determine the best model order for the channel model.

Once the data analysis unit 49 has determined the best model orders (k and r), it then analyses the samples generated by the Gibbs sampler 41 which are stored in the store 53 of the RAM 47-2, in order to determine parameter values that are most representative of those samples.

It does this by determining a histogram for each of the parameters from which it determines the most representative parameter value. To generate the histogram, the data analysis unit 49 determines the maximum and minimum sample value which was drawn by the Gibbs sampler and then divides the range of parameter values between this minimum and maximum value into a predetermined number of sub-ranges or bins. The data analysis unit 49 then assigns each of the sample values into the appropriate bins and counts how many samples are allocated to each bin. It then uses these counts to

calculate a weighted average of the samples (with the weighting used for each sample depending on the count for the corresponding bin), to determine the most representative parameter value (known as the minimum mean square estimate (MMSE)). Figure 9b illustrates an example histogram which is generated for the variance (σ_e^2) of the process noise, from which the data analysis unit 49 determines that the variance representative of the sample is 0.3149.

10

In determining the AR filter coefficients (a_i for $i = 1$ to k), the data analysis unit 49 determines and analyses a histogram of the samples for each coefficient independently. Figure 9c shows an exemplary histogram obtained for the third AR filter coefficient (a_3), from which the data analysis unit 49 determines that the coefficient representative of the samples is -0.4977.

15

In this embodiment, the data analysis unit 49 only outputs the AR filter coefficients which are passed to the channel encoder 71 shown in Figure 1. The AR filter coefficients (and the remaining parameter values determined by the data analysis unit 49) are also stored in the RAM 47-2 for use during the processing of the next frame of speech.

20

25

As the skilled reader will appreciate, a speech processing technique has been described above which uses statistical analysis techniques to determine sets of AR filter coefficients representative of an input speech signal. The technique is more robust and accurate than prior art techniques which employ maximum likelihood estimators to determine the AR filter coefficients. This is because the statistical analysis of each frame uses knowledge obtained from the processing of the previous frame.

In addition, with the analysis performed above, the model order for the AR filter model is not assumed to be constant and can vary from frame to frame. In this way, the optimum number of AR filter coefficients can be used to represent the speech within each frame. As a result, the AR filter coefficients output by the statistical analysis unit 21 will more accurately represent the corresponding input speech. In contrast, with the prior art linear prediction systems, the number of AR coefficients is assumed to be constant and hence these prior art techniques tend to over parametrise the speech in order to ensure that information is not lost. As a result, with the statistical analysis described above, the amount of data which has to be transmitted from the

transmitter to the receiver will be less than with the prior art systems which assume a fixed size of AR filter model.

5 Further still, since the underlying process model that is used separates the speech source from the channel, the AR filter coefficients that are determined will be more representative of the actual speech and will be less likely to include distortive effects of the channel.

10

Further still, since variance information is available for each of the parameters, this provides an indication of the confidence of each of the parameter estimates. This is in contrast to maximum likelihood and least
15 squares approaches, such as linear prediction analysis, where point estimates of the parameter values are determined.

Alternative Embodiments

20 In the above embodiment, the statistical analysis unit was used in order to parameterise the input speech signal for onward transmission to a remote receiver. It also generated a number of other parameter values (such as the process noise variances and the channel model
25 coefficients), but these were not output by the

statistical analysis unit. As those skilled in the art will appreciate, the AR coefficients and some of the other parameters which are calculated by the statistical analysis unit can also be used in the transmission system. For example, the variance information of the AR coefficients can be used to control the type of channel encoding which is employed by the channel encoder 71, since this variance information is indicative of the quality of the input speech signal. For example, different channel encoding techniques may be used in dependence upon the quality of the input speech signal, with the particular technique being used depending upon the quality of the speech within each frame.

In the above embodiment, the speech generation process was modelled as an auto-regressive (AR) process and the channel was modelled as a moving average (MA) process. As those skilled in the art will appreciate, other signal models may be used. However, these models are preferred because it has been found that they suitably represent the speech source and the channel they are intended to model.

In the above embodiment, a speech quality indicator was determined for each frame of speech that is processed and

the channel encoder 71 encodes each frame of speech in dependence upon the quality measure determined for each frame. As those skilled in the art will appreciate, this is not essential. For example, a moving average quality indicator may be determined and the channel encoder may be arranged to change its encoding technique when the moving average goes above or below a predetermined threshold value.

10 In the above embodiments, Gaussian and Inverse Gamma distributions were used to model the various prior probability density functions of equation (19). As those skilled in the art of statistical analysis will appreciate, the reason these distributions were chosen is
15 that they are conjugate to one another. This means that each of the conditional probability density functions which are used in the Gibbs sampler will also either be Gaussian or Inverse Gamma. This therefore simplifies the task of drawing samples from the conditional probability
20 densities. However, this is not essential. The noise probability density functions could be modelled by Laplacian or student-t distributions rather than Gaussian distributions. Similarly, the probability density functions for the variances may be modelled by a
25 distribution other than the Inverse Gamma distribution.

For example, they can be modelled by a Rayleigh distribution or some other distribution which is always positive. However, the use of probability density functions that are not conjugate will result in increased complexity in drawing samples from the conditional densities by the Gibbs sampler.

Additionally, whilst the Gibbs sampler was used to draw samples from the probability density function given in equation (19), other sampling algorithms could be used. For example the Metropolis-Hastings algorithm (which is reviewed together with other techniques in a paper entitled "Probabilistic inference using Markov chain Monte Carlo methods" by R. Neal, Technical Report CRG-TR-93-1, Department of Computer Science, University of Toronto, 1993) may be used to sample this probability density.

In the above embodiment, a Simulation Smoother was used to generate estimates for the raw speech samples. This Simulation Smoother included a Kalman filter stage and a smoothing filter stage in order to generate the estimates of the raw speech samples. In an alternative embodiment, the smoothing filter stage may be omitted, since the Kalman filter stage generates estimates of the raw speech

(see equation (33)). However, these raw speech samples were ignored, since the speech samples generated by the smoothing filter are considered to be more accurate and robust. This is because the Kalman filter essentially generates a point estimate of the speech samples from the joint probability density function $p(\underline{s}(n) | \underline{a}, k, \sigma_e^2)$, whereas the Simulation Smoother draws a sample from this probability density function.

10 In the above embodiment, a Simulation Smoother was used in order to generate estimates of the raw speech samples. It is possible to avoid having to estimate the raw speech samples by treating them as "nuisance parameters" and integrating them out of equation (19). However, this is
15 not preferred, since the resulting integral will have a much more complex form than the Gaussian and Inverse Gamma mixture defined in equation (19). This in turn will result in more complex conditional probabilities corresponding to equations (20) to (30). In a similar
20 way, the other nuisance parameters (such as the coefficient variances or any of the Inverse Gamma, alpha and beta parameters) may be integrated out as well. However, again this is not preferred, since it increases the complexity of the density function to be sampled
25 using the Gibbs sampler. The technique of integrating

out nuisance parameters is well known in the field of statistical analysis and will not be described further here.

5 In the above embodiment, the data analysis unit analysed the samples drawn by the Gibbs sampler by determining a histogram for each of the model parameters and then determining the value of the model parameter using a weighted average of the samples drawn by the Gibbs
10 sampler with the weighting being dependent upon the number of samples in the corresponding bin. In an alternative embodiment, the value of the model parameter may be determined from the histogram as being the value of the model parameter having the highest count.
15 Alternatively, a predetermined curve (such as a bell curve) could be fitted to the histogram in order to identify the maximum which best fits the histogram.

In the above embodiment, the statistical analysis unit
20 modelled the underlying speech production process with a separate speech source model (AR filter) and a channel model. Whilst this is the preferred model structure, the underlying speech production process may be modelled without the channel model. In this case, there is no
25 need to estimate the values of the raw speech samples

using a Kalman filter or the like, although this can still be done. However, such a model of the underlying speech production process is not preferred, since the speech model will inevitably represent aspects of the channel as well as the speech. Further, although the statistical analysis unit described above ran a model order selection routine in order to allow the model orders of the AR filter model and the channel model to vary, this is not essential. In particular, the model order of the AR filter model and the channel model may be fixed in advance, although this is not preferred since it will inevitably introduce errors into the representation.

In the above embodiments, the speech that was processed was received from a user via a microphone. As those skilled in the art will appreciate, the speech may be received from a telephone line or may have been stored on a recording medium. In this case, the channel model will compensate for this so that the AR filter coefficients representative of the actual speech that has been spoken should not be significantly affected.

In the above embodiments, during the running of the model order selection routine, a new model order was proposed by drawing a random variable from a predetermined

Laplacian distribution function. As those skilled in the art will appreciate, other techniques may be used. For example the new model order may be proposed in a deterministic way (ie under predetermined rules), provided that the model order space is sufficiently sampled.

CLAIMS

1. An audio encoding system comprising:

5 a memory for storing a predetermined function which gives, for a given set of audio signal values, a probability density for parameters of a predetermined audio model which is assumed to have generated the set of audio signal values, the probability density defining, for a given set of model parameter values, the probability that the predetermined audio model has those parameter values, given that the model is assumed to have generated the set of audio signal values;

10 means for receiving a set of audio signal values representative of an input audio signal;

15 means for applying the set of received audio signal values to said stored function to give the probability density for said model parameters for the set of received audio signal values;

20 means for processing said function with said set of received audio signal values applied, to derive samples of parameter values from said probability density;

25 means for analysing at least some of said derived samples of parameter values to determine parameter values that are representative of the set of received audio signal values; and

means for encoding said determined parameter values to generate encoded data representative of the received audio signal values.

- 5 2. A system according to claim 1, wherein said processing means is operable to draw samples iteratively from said probability density function.
- 10 3. A system according to claim 2, wherein said processing means comprises a Gibbs sampler.
- 15 4. A system according to claim 2 or 3, wherein said analysing means is operable to determine a histogram of said drawn samples and wherein said values of said parameters are determined from said histogram.
- 20 5. A system according to claim 4, wherein said processing means is operable to determine said values of said first parameters using a weighted sum of said drawn samples and wherein the weighting for each sample is determined from said histogram.
- 25 6. A system according to any preceding claim, wherein said receiving means is operable to receive a sequence of sets of signal values representative of an input audio

signal and wherein said applying means, processing means and analysing means are operable to perform their function with respect to each set of received audio signal values to determine parameter values that are representative of each set of received audio signal values.

5
7. A system according to claim 6, wherein said processing means is operable to use the values of parameters obtained during the processing of a preceding set of signal values as initial estimates for the values of the corresponding parameters for a current set of signal values being processed.

10
8. A system according to claim 6 or 7, wherein said sets of signal values in said sequence are non-overlapping.

15
9. A system according to any of claims 6 to 8, wherein said processing means comprises means for varying the number of parameters used to represent the audio signal within each set of audio signal values.

20
10. A system according to any preceding claim, wherein said audio model comprises an auto-regressive process

model and wherein said parameters include auto-regressive model coefficients.

5 11. A system according to any preceding claim, wherein said received set of audio signal values are representative of an input speech signal.

10 12. A system according to claim 11, wherein said received set of speech signal values are representative of a speech signal generated by a speech source as distorted by a transmission channel between the speech source and the receiving means; wherein said predetermined function includes a first part having first parameters which models said source and a second part
15 having second parameters which models said channel; wherein said processing means is operable to derive samples of at least said first parameters; and wherein said analysing means is operable to determine values of said first parameters that are representative of said
20 speech generated by said speech source before it was distorted by said transmission channel.

25 13. A system according to claim 12, wherein said function is in terms of a set of raw speech signal values representative of speech generated by said source before

being distorted by said transmission channel, wherein the system further comprises second processing means for processing the received set of signal values with initial estimates of said first and second parameters, to
5 generate an estimate of the raw speech signal values corresponding to the received set of signal values and wherein said applying means is operable to apply said estimated set of raw speech signal values to said function in addition to said set of received signal
10 values.

14. A system according to claim 13, wherein said second processing means comprises a simulation smoother.

15 15. A system according to claim 13 or 14, wherein said second processing means comprises a Kalman filter.

16. A system according to any of claims 12 to 15, wherein said second part is a moving average model and
20 said second parameters comprise moving average model coefficients.

17. A system according to any preceding claim, further comprising means for evaluating said probability density
25 function for the set of received signal values using one

or more of said drawn samples of parameter values for different numbers of parameter values, to determine respective probabilities that the predetermined signal model has those parameter values and wherein said processing means is operable to process at least some of said drawn samples of parameter values and said evaluated probabilities to determine said values of said parameters that are representative of the received audio signal.

10 18. A system according to any preceding claim, wherein said analysing means is operable to analyse at least some of said derived samples of parameter values to determine a measure of the variance of at least some of said samples of parameter values; wherein said system further
15 comprises means for determining an indication of the quality of the received audio signal using said variance measure; and wherein said encoding means is operable to encode said determined parameter values in dependence upon the determined quality indication.

20 19. A system according to claim 18, wherein said encoding means is operable to encode said parameter values using a first encoding technique if said quality indication is above a predetermined value and is operable
25 to encode said parameter values using a second encoding

technique if said quality indication is below said value.

20. A system according to claim 19, wherein said first
5 encoding technique is operable to minimise the data to be
transmitted and wherein said second encoding technique
is operable to minimise information lost in the encoding.

21. An audio transmission system comprising:
a transmission unit comprising: means for receiving
10 an audio signal; an audio encoding system according to
any preceding claim for generating encoded parameter
values representative of received audio signal values;
and means for transmitting the encoded parameter values;
and a receiver unit comprising means for receiving the
15 transmitted parameter values; and means for processing
the received parameter values to generate an output
signal in dependence thereon.

22. A system according to claim 21, wherein said
20 processing means of said receiving unit comprises speech
synthesis means for generating a synthesised speech
signal in dependence upon the received parameter values.

23. A system according to claim 21 or 22, wherein said
25 processing means of said receiving unit comprises a

speech recognition system which operable to compare the received parameter values with stored reference models to generate a recognition result.

5 24. A system according to claim 21 when dependent upon claim 19, wherein said transmission unit is operable to transmit said quality indication to said receiving unit and wherein said receiving unit is operable to receive said quality indication and to decode said encoded.
10 parameters in dependence upon the received quality indication.

25. A system according to claim 24, wherein said receiving unit is operable to decode said encoded
15 parameter values in accordance with a first decoding technique if said quality indication has a value above a predetermined threshold value and is operable to decode said encoded parameter values in accordance with a second decoding technique if said quality indication is below
20 said predetermined value.

26. An audio transmission system comprising a transmitter and receiver, wherein

the transmitter comprises:

25 means for receiving an input audio signal;

means for determining a measure of the quality of the input audio signal;

means for encoding data representative of the audio signal in dependence upon the determined quality measure; and

means for transmitting the encoded audio data; and wherein

said receiver comprises:

means for receiving the encoded audio data;

means for decoding the transmitted audio data;

and

means for outputting the decoded audio data.

27. A system according to claim 26, wherein said transmitter is operable to transmit said quality measure and wherein said decoder is operable to decode said encoded audio data in dependence upon the received quality measure.

28. A system according to claim 26 or 27, wherein said encoder is operable to encode said audio signal in accordance with a first encoding technique if said quality measure is above a predetermined threshold and is operable to encode said audio signal in accordance with a second encoding technique if said quality measure is

below a predetermined threshold.

5 29. A system according to claim 28, wherein said first encoding technique is operable to minimise the data to be transmitted and where said second encoding technique is operable to minimise information lost in the encoding.

30. An audio transmission system comprising a transmission unit and a receiving unit, wherein
10 the transmission unit comprises:

15 a memory for storing a predetermined function which gives, for a given set of audio signal values, a probability density for parameters of a predetermined audio model which is assumed to have generated the set of audio signal values, the probability density defining, for a given set of model parameter values, the probability that the predetermined audio model has those parameter values, given that the model is assumed to have generated the set of audio signal values;

20 means for receiving a set of audio signal values representative of an input audio signal;

25 means for applying the set of received audio signal values to said stored function to give the probability density for said model parameters for the set of received audio signal values;

means for processing said function with said set of received audio signal values applied, to derive samples of parameter values from said probability density;

5 means for analysing at least some of said derived samples of parameter values to determine parameter values that are representative of the set of received audio signal values; and

means for transmitting said determined
10 parameter values; and wherein

the receiver unit comprises:

means for receiving said transmitted parameter values; and

means for processing the received parameter
15 values to generate an output signal in dependence thereon.

31. A system according to claim 30, wherein said transmission unit further comprises means for encoding
20 the determined parameter values and wherein said receiving unit comprises means for decoding the encoded parameter values.

32. A transmitter apparatus comprising:
25 a memory for storing a predetermined function which

gives, for a given set of audio signal values, a probability density for parameters of a predetermined audio model which is assumed to have generated the set of audio signal values, the probability density defining, for a given set of model parameter values, the probability that the predetermined audio model has those parameter values, given that the model is assumed to have generated the set of audio signal values;

means for receiving a set of audio signal values representative of an input audio signal;

means for applying the set of received audio signal values to said stored function to give the probability density for said model parameters for the set of received audio signal values;

means for processing said function with said set of received audio signal values applied, to derive samples of parameter values from said probability density;

means for analysing at least some of said derived samples of parameter values to determine parameter values that are representative of the set of received audio signal values; and

means for transmitting said determined parameter values.

33. A transmitter apparatus comprising:

means for receiving an input audio signal;

means for determining a measure of the quality of the input audio signal;

means for encoding data representative of the audio signal in dependence upon the determined quality measure; and

means for transmitting the encoded data.

34. A transmitter according to claim 33, wherein said transmitting means is operable to transmit said quality measure in addition to said encoded data.

35. A transmitter according to claim 33 or 34, wherein said encoder is operable to encode said audio signal in accordance with a first encoding technique if said quality measure is above a predetermined threshold and is operable to encode said audio signal in accordance with a second encoding technique if said quality measure is below a predetermined threshold.

36. A transmitter according to claim 35, wherein said first encoding technique is operable to minimise the data to be transmitted and wherein said second encoding technique is operable to minimise information lost in the encoding.

37. An audio encoding method comprising the steps of:

storing a predetermined function which gives, for a given set of audio signal values, a probability density for parameters of a predetermined audio model which is assumed to have generated the set of audio signal values, the probability density defining, for a given set of model parameter values, the probability that the predetermined audio model has those parameter values, given that the model is assumed to have generated the set of audio signal values;

receiving a set of audio signal values representative of an input audio signal at a receiver;

applying the set of received audio signal values to said stored function to give the probability density for said model parameters for the set of received audio signal values;

processing said function with said set of received audio signal values applied, to derive samples of parameter values from said probability density;

analysing at least some of said derived samples of parameter values to determine parameter values that are representative of the set of received audio signal values; and

encoding said determined parameter values to

generate encoded data representative of the received audio signal values.

5

38. A method according to claim 37, wherein said processing step draws samples iteratively from said probability density function.

10

39. A method according to claim 38, wherein said processing step uses a Gibbs sampler.

15

40. A method according to claim 38 or 39, wherein said analysing step determines a histogram of said drawn samples and wherein said values of said parameters are determined from said histogram.

20

41. A method according to claim 40, wherein said processing step determines said values of said first parameters using a weighted sum of said drawn samples and wherein the weighting for each sample is determined from said histogram.

25

42. A method according to any of claims 37 to 41, wherein said receiving step receives a sequence of sets of signal values representative of an input audio signal and wherein said applying step, processing step and

analysing step are performed for each set of received audio signal values to determine parameter values that are representative of each set of received audio signal values.

5

10

43. A method according to claim 42, wherein said processing step uses the values of parameters obtained during the processing of a preceding set of signal values as initial estimates for the values of the corresponding parameters for a current set of signal values being processed.

15

44. A method according to claim 42 or 43, wherein said sets of signal values in said sequence are non-overlapping.

20

45. A method according to any of claims 42 to 44, wherein said processing step comprises the step of varying the number of parameters used to represent the audio signal within each set of audio signal values.

25

46. A method according to any of claims 37 to 45, wherein said audio model comprises an auto-regressive process model and wherein said parameters include auto-regressive model coefficients.

47. A method according to any of claims 37 to 46, wherein said received set of audio signal values are representative of an input speech signal.

5 48. A method according to claim 47, wherein said received set of speech signal values are representative of a speech signal generated by a speech source as distorted by a transmission channel between the speech source and the receiver; wherein said predetermined
10 function includes a first part having first parameters which models said source and a second part having second parameters which models said channel; wherein said processing step derives samples of at least said first parameters; and wherein said analysing step determines
15 values of said first parameters that are representative of said speech generated by said speech source before it was distorted by said transmission channel.

20 49. A method according to claim 48, wherein said function is in terms of a set of raw speech signal values representative of speech generated by said source before being distorted by said transmission channel, further comprising a second processing step for processing the received set of signal values with initial estimates of
25 said first and second parameters, to generate an estimate

of the raw speech signal values corresponding to the received set of signal values and wherein said applying step applies said estimated set of raw speech signal values to said function in addition to said set of received signal values.

5

50. A method according to claim 49, wherein said second processing step uses a simulation smoother.

10

51. A method according to claim 49 or 50, wherein said second processing step uses a Kalman filter.

15

52. A method according to any of claims 48 to 50 wherein said second part is a moving average model and said second parameters comprise moving average model coefficients.

20

53. A method according to any of claims 37 to 52, further comprising the step of evaluating said probability density function for the set of received signal values using one or more of said drawn samples of parameter values for different numbers of parameter values, to determine respective probabilities that the predetermined signal model has those parameter values and

25

wherein said processing step processes at least some of

said drawn samples of parameter values and said evaluated probabilities to determine said values of said parameters that are representative of the received audio signal.

5 54. A method according to any of claims 37 to 53,
wherein said analysing step analyses at least some of
said derived samples of parameter values to determine a
measure of the variance of at least some of said samples
of parameter values; further comprising the step of
10 determining an indication of the quality of the received
audio signal; and wherein said encoding step encodes
said determined parameter values in dependence upon the
determined quality indication.

15 55. A method according to claim 54, wherein said
encoding step encodes said parameter values using a first
encoding technique if said quality indication is above a
predetermined value and encodes said parameter values
using a second encoding technique if said quality
20 indication is below said value.

56. A method according to claim 55, wherein said first
encoding technique is operable to minimise the data to be
transmitted and wherein said second encoding technique
25 is operable to minimise information lost in the

encoding.

57. An audio transmission method comprising the steps of:

- 5 receiving an audio signal at a transmission unit;
 encoding the audio signal using a method according
to any of claims 37 to 56 to generate encoded parameter
values representative of the audio signal; and
 transmitting the encoded parameters values;
10 receiving the transmitted encoded parameter values
at a receiver unit;
 decoding the received encoded parameter values;
 processing the decoded parameter values to generate
and output signal in dependence thereon.

15

58. A method according to any of claims 37 to 53,
wherein said processing step at said receiving unit
comprises speech synthesis means for generating a
synthesised speech signal in dependence upon the received
20 parameter values.

20

59. A method according to any of claims 37 to 54,
wherein said processing step at said receiving unit uses
a speech recognition system to compare the received
25 parameter values with stored reference models and to

generate a recognition result.

5 60. A method according to claim 57 when dependent upon claim 54, further comprising the step of transmitting said quality indication to said receiving unit and, at said receiving unit, the steps of receiving said quality indication and decoding said encoded parameters in dependence upon the received quality indication.

10 61. A method according to claim 60, comprising the step of, at said receiving unit, decoding said encoded parameter values in accordance with a first decoding technique if said quality indication has a value above a predetermined threshold value and decoding said encoded
15 parameter values in accordance with a second decoding technique if said quality indication is below said predetermined value.

20 62. An audio transmission method using a transmitter and receiver, the method comprising the steps of:

at the transmitter:

receiving an audio signal;

determining a measure of the quality of the received audio signal;

25 encoding data representative of the audio

signal in dependence upon the determined quality measure; and

transmitting the encoded data; and

at said receiver:

5 receiving the encoded audio data;
decoding the transmitted audio data; and
outputting the decoded audio data.

63. A method according to claim 62, further comprising
10 the step of, at said transmitter, transmitting said quality measure and wherein said decoding step decodes said encoded audio data in dependence upon the received quality measure.

15 64. A method according to claim 61 or 62, wherein said encoding step encodes said audio signal in accordance with a first encoding technique if said quality measure is above a predetermined threshold and encodes said audio signal in accordance with a second encoding technique if
20 said quality measure is below a predetermined threshold.

65. A method according to claim 63, wherein said first encoding technique is operable to minimise the data to be transmitted and where said second encoding technique is
25 operable to minimise information lost in the encoding.

66. An audio transmission method using a transmission unit and a receiving unit, wherein the method comprises the steps of:

5 at the transmission unit:

 storing a predetermined function which gives, for a given set of audio signal values, a probability density for parameters of a predetermined audio model which is assumed to have generated the set of audio signal values, the probability density defining, for a given set of model parameter values, the probability that the predetermined audio model has those parameter values, given that the model is assumed to have generated the set of audio signal values;

15 receiving a set of audio signal values representative of an input audio signal;

 applying the set of received audio signal values to said stored function to give the probability density for said model parameters for the set of received audio signal values;

20 processing said function with said set of received audio signal values applied, to derive samples of parameter values from said probability density;

 analysing at least some of said derived samples of parameter values to determine parameter values that

25

are representative of the set of received audio signal values; and

transmitting said determined parameter values;

and

5 at the receiver unit:

receiving said transmitted parameter values;

and

processing the received parameter values to generate an output signal in dependence thereon.

10

67. A computer readable medium storing computer executable process steps to cause a programmable computer apparatus to perform the method of any of claims 37 to 66.

15

68. Processor implementable process steps for causing a programmable computing device to perform the method according to any of claims 37 to 66.


20



ABSTRACTSPEECH PROCESSING SYSTEM

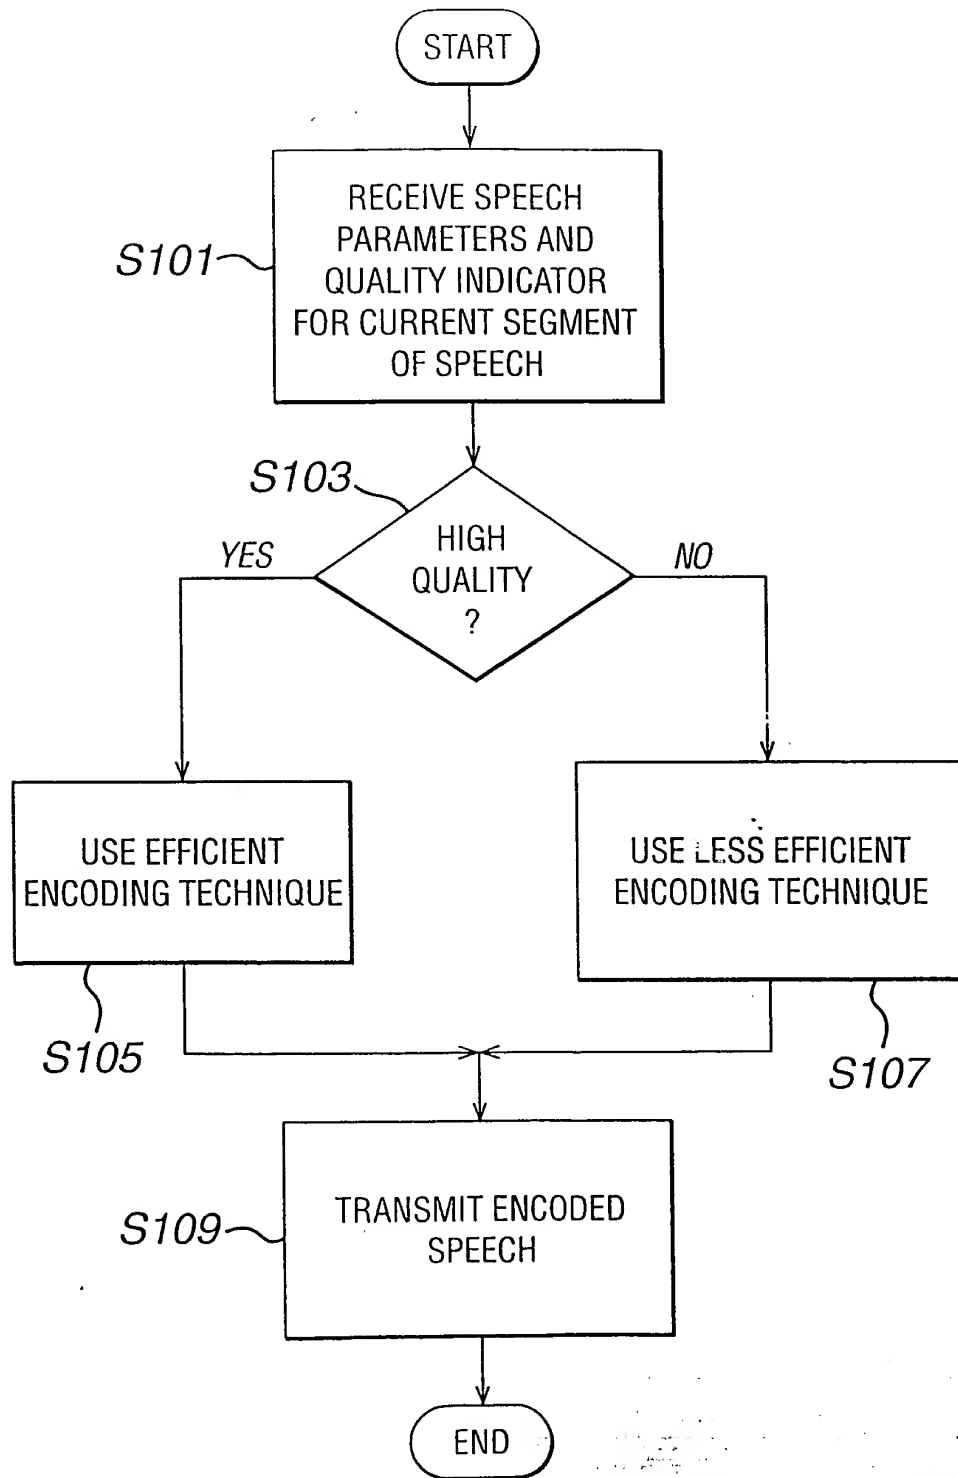
5 An audio transmission system is provided which is
operable to receive sets of signals representative of a
speech signal generated by a speech source. The system
is operable to determine parameters representative of
each of the sets of signal values, to encode these
parameter values and then to transmit these parameter
10 values to a receiver circuit. At the receiver, the
encoded parameters are decoded to recover the parameter
values. The recovered parameter values may then be used
to resynthesise the speech using a speech synthesiser or
they may be used by a speech recognition system to
15 generate a recognition result.

FIG. 1



2/11

FIG. 2a



3/11

FIG. 2b





FIG. 3



5/11

FIG. 4

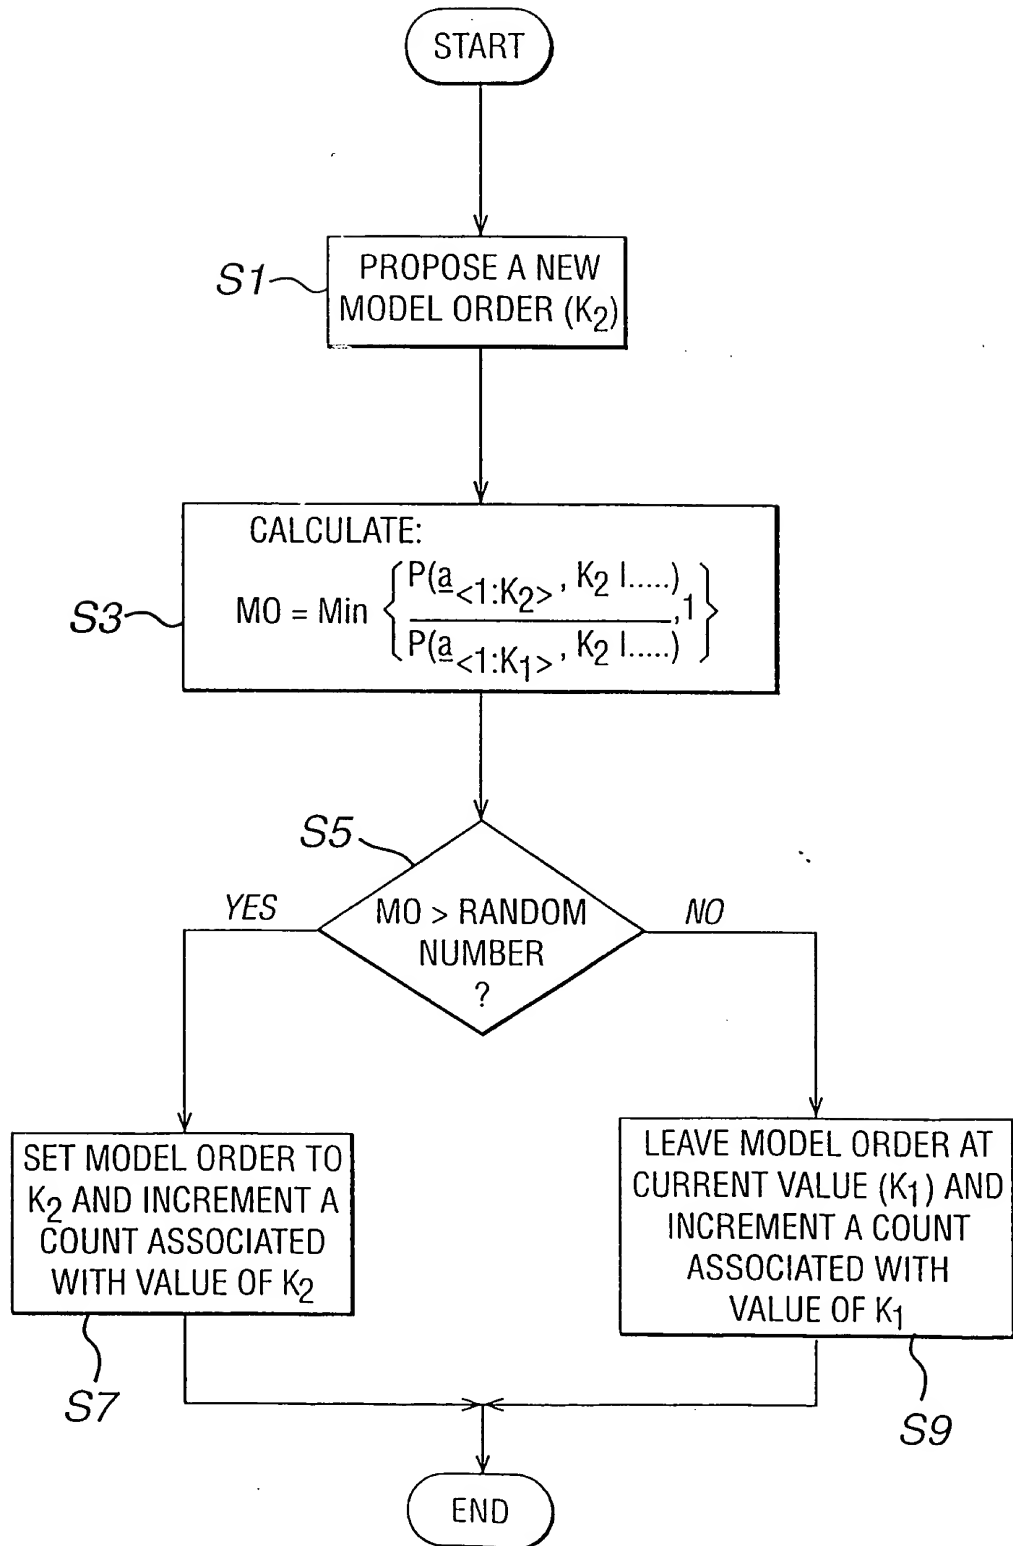


FIG. 5

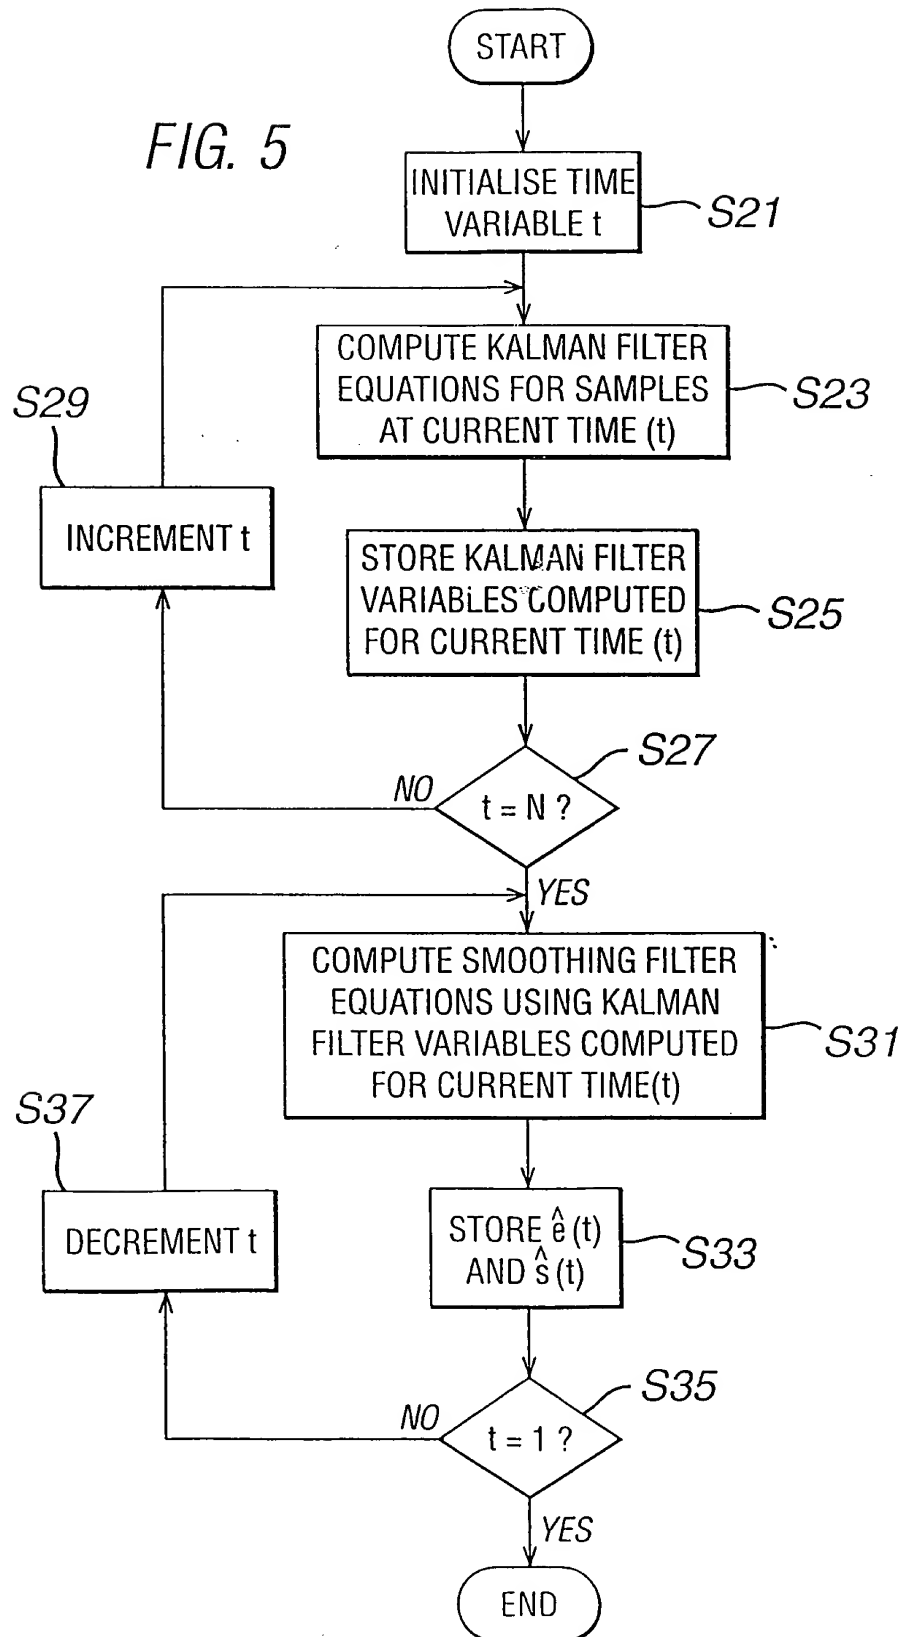


FIG. 6



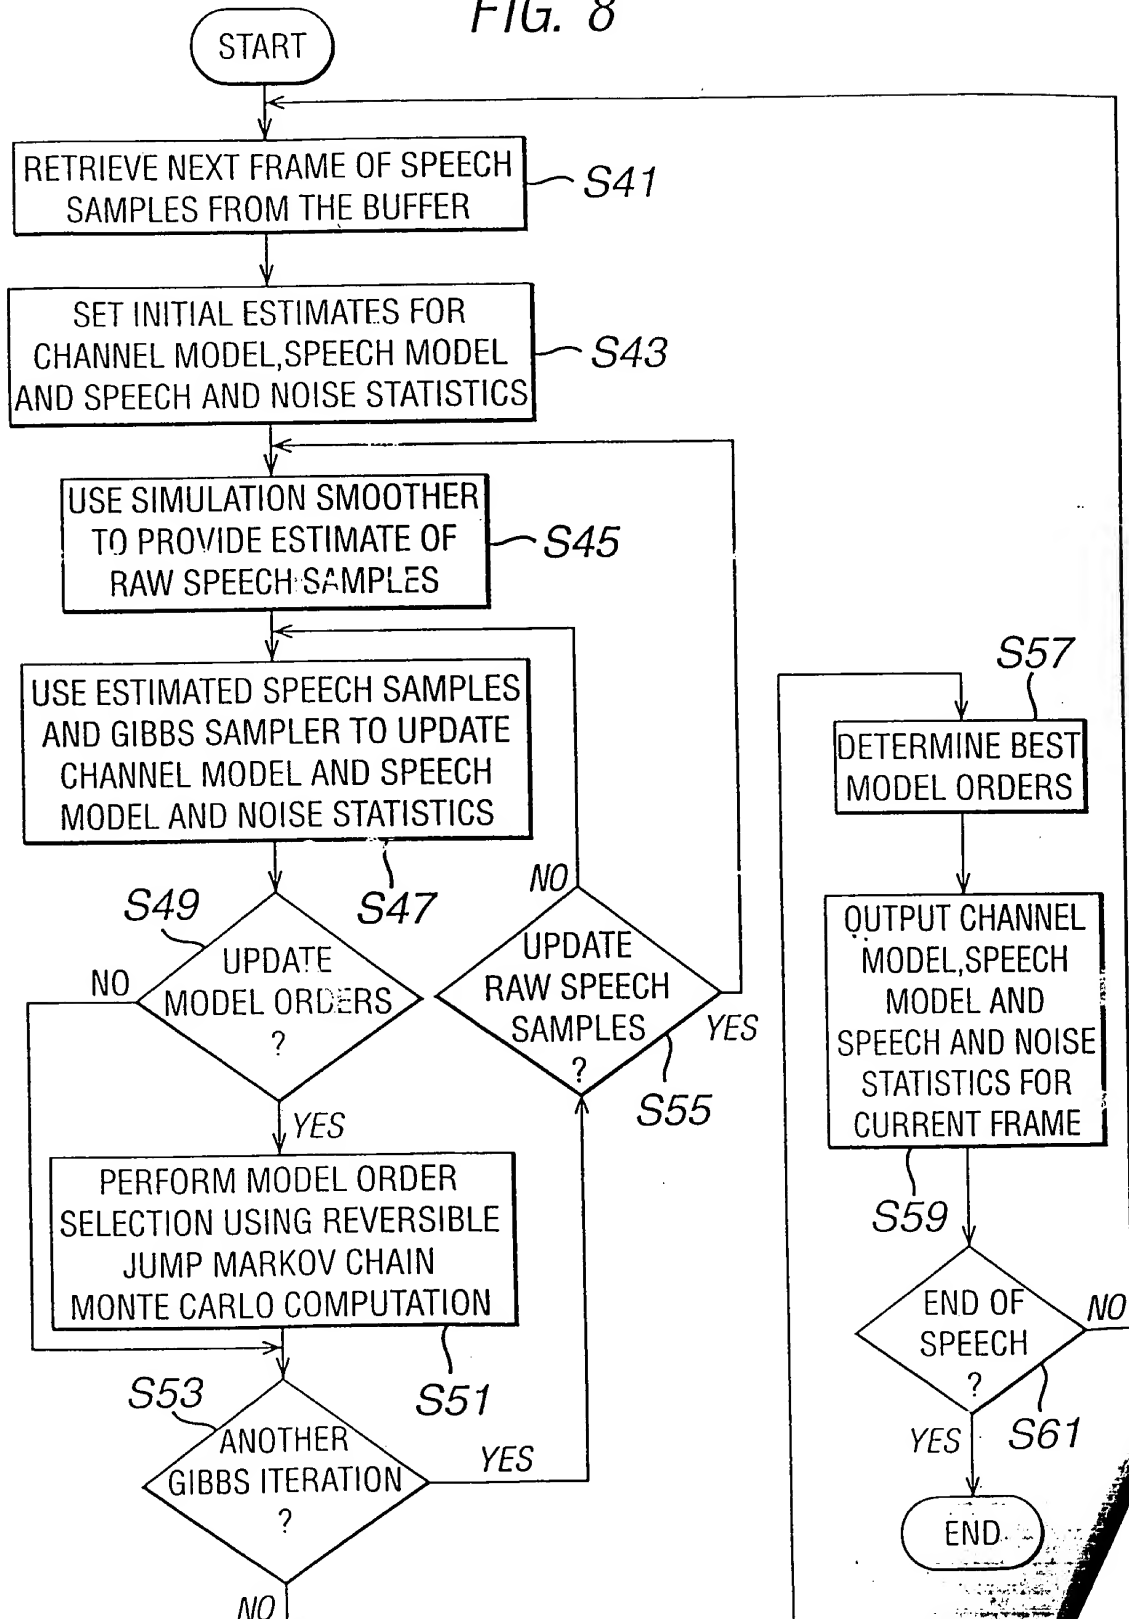


FIG. 7



9/11

FIG. 8



10/11

FIG. 9a





FIG. 9b



11/11

FIG. 9c

